

# CISE NREN FCCSET reports, papers, and interviews

Catalog number: 102740244  
Gordon Bell NSF-CISE Records  
Virtual Bankers Box<sup>1</sup> 2



---

<sup>1</sup> To view individual document titles within this PDF, set your PDF reader to display bookmarks.

Contents: Bell-NSF-CISE archives in three VBBs (Virtual Banker Boxes) of approx. 750 pages. The boxes contain budgets, goals, organizations, presentations, and correspondence. Box 2 is about NSF's role in creating the National Research and Education Network (NREN) aka the Internet. Box 2 contains a number of interviews by persons creating the history of CISE and the Internet.	MB	Bookmarked Items	Pages
<u>Bell CISE Directorate of NSF. Charter, Budgets, Goals, Org. 1986-1987 VBB1.</u> Contains rationale for a new NSF Division taken from Engineering, Math and Science, and Social Science. Talk to Congress, CS community, including CRA.	3.6	26	100
<u>Bell CISE NREN FCCSET reports, papers, &amp; Aspray, IEEE, Kleinrock, van Houweling interviews VBB2.</u> This folder is about NSF's role in creating the National Research and Education Network (NREN) aka the Internet and has my interviews with people and organizations writing the history of CISE and the Internet	18.2	16	441
<u>Bell CISE NSF Centers for ASC division VBB3.</u> NSF Supercomputer Centers that were established at this time. The funding and operation was somewhat contentious since the computer science community had no interest in scientific computing in 1986. Thirty years later the condition is about the same even though Computational Science is a CS sub-discipline.	58.9	47	218
	80.7	89	759

A Report to the Office of  
Science and Technology Policy on  
Computer Networks to Support Research  
In the United States

\*\*\*\*\*

A Study of Critical Problems and Future Options

Volume I  
Recommendations

November 1987

Document prepared by the  
Computing and Communication Division,  
Los Alamos National Laboratory, Los Alamos, New Mexico.\*

This report was prepared as an account of work sponsored by agencies of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

\*The Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

## CONTENTS

---

<b>OVERVIEW</b> .....	1
<b>TRANSMITTAL LETTER</b> .....	2
<b>1. EXECUTIVE SUMMARY</b> .....	3
<b>2. COMPUTER NETWORK STUDY</b> .....	4
2.1. Background .....	4
2.2. Findings .....	5
2.3. Requirements .....	5
2.4. Recommendations .....	6
2.5. Benefits .....	8
2.6. Action Plan .....	9
2.7. Stage 1 The Internet .....	9
2.8. Stage 2 The National Research Network .....	10
2.9. Stage 3 The Enhanced National Research Network .....	10
<b>APPENDIX A - PUBLIC LAW 99-383</b> .....	12
<b>APPENDIX B - FCCSET COMMITTEE MEMORANDUM</b> .....	13
Background .....	14
Purpose .....	14
Implementation .....	14
Membership .....	14
Liaison Participants .....	15
Executive (Steering) Committee .....	15
Administrative Provisions .....	15
Duration .....	15
Compensation .....	15
Determination .....	15
FCCSET Committee on Computer Research and Applications .....	16
<b>APPENDIX C - INDEX TO OTHER VOLUMES</b> .....	17
<b>APPENDIX D - NETWORKING REQUIREMENTS FOR MOLECULAR BIOLOGY</b> .....	20



## OVERVIEW

---

This Computer Network Study is published in three volumes. Volume I contains the FCCSET recommendations to the Office of Science and Technology Policy on developing computer networks to support research in the U.S. Volume II contains the summaries of the February 1987 workshop discussions, which focused on six topics: access requirements and future alternatives; special requirements for supercomputer networks; internetwork concepts; future standards and services requirements; security issues; and the government role in networking. Volume III contains white papers that the Network Study Group invited on networking trends, requirements, concepts, applications, and plans.

### Computer Network Study Group

F. Ron Bailey, National Aeronautics and Space Administration  
Gordon Bell, National Science Foundation  
James Burrows, National Bureau of Standards  
John Cavallini, Department of Energy  
Michael Corrigan, Office of the Secretary of Defense  
Paul Huray, Office of Science and Technology Policy  
Thomas Kitchens, Department of Energy  
James Oberthaler, Department of Health and Human Services  
Dennis Perry, Defense Advanced Research Projects Agency  
Shirley Radack, National Bureau of Standards  
Rudi Saenger, Naval Research Laboratory  
Daniel VanBelleghem, National Science Foundation  
Stephen Wolff, National Science Foundation

## TRANSMITTAL LETTER

---

Dr. Paul Huray  
Chairman, Executive (Steering) Committee  
Committee on Computer Research and Applications  
Federal Coordinating Council on Science, Engineering and Technology

Dear Dr. Huray:

I am pleased to transmit to you this Computer Network Study which was done by the Federal Coordinating Council for Science, Engineering and Technology at the request of the Office of Science and Technology Policy. This study responds to a charge of the 99th Congress for "a study of critical problems and current and future options regarding communications networks for research computers, including supercomputers, at universities and Federal research facilities in the United States." (Public Law 99-383, August 21, 1986). The Congressional charge asked that requirements for supercomputers be addressed within one year and requirements for all research computers be addressed within two years. Requirements for both supercomputers and research computers are addressed in this report; therefore, no second year study is planned.

Our principal recommendation is that an advanced computer network be designed and developed to interconnect academic, industrial, and government research facilities in the U.S. This proposed network offers a challenging opportunity to enhance the research capabilities throughout this country and to improve the networking capabilities of U.S. industry. To support this innovative project, a vigorous and focused program of research and development is needed, starting immediately and continuing for a 15 year period, during which time the network will be developed.

In conducting this study, we called upon the help of many experts from government, industry, and academia. White papers were invited on networking trends, requirements, concepts, applications, and plans. A workshop involving nearly 100 researchers, network users, network suppliers, and policy officials was held to air ideas, gather information and develop the foundation for our recommendations.

I believe that this study points the way to future progress in many areas of high technology research in the U.S., and I thank all of the people who have contributed -- the workshop participants; the chairs of the workshop groups; the San Diego Supercomputer Center which hosted the workshop; authors of the white papers; the Department of Energy and the Los Alamos National Laboratory staff who edited and published the 3 volumes of the report; and the members of the FCCSET group that conducted the study.

Sincerely,

Gordon Bell  
Chairman

Subcommittee on Computer Networking,  
Infrastructure, and Digital Communications

Committee on Computer Research and Applications,  
Federal Coordinating Council on Science,  
Engineering and Technology



## 1. EXECUTIVE SUMMARY

---

A strong national effort, supported by the Federal government, is needed to improve computer networks in the U.S. and to improve the access of U.S. researchers to computing and research facilities. Today's technology is not adequate to support access to high performance computing or requirements for researchers to collaborate through computer networks. Over the next 15 years, there will be a need for a 100,000 times increase in national network capacity to enable researchers to exploit computer capabilities for representing complex data in visual form, for manipulating and interacting with this complex data and for sharing large data bases with other researchers.

The key to improving the ability of computers to serve U.S. science is better coordinated efforts of agencies that support research networks, and a new initiative to carry out engineering and research in improved data communications speeds, switching technology, network security, and interoperability standards. Rough estimates of the costs of carrying out this work are included with this study.

A plan of action is recommended to conduct a three stage program starting with the internetworking and upgrading of current agency networks and progressing to higher speed data communication services reaching virtually every university and industry research facility in the U.S.

- As the first step, the current Internet system developed by the Defense Advanced Research Projects Agency and the networks supported by agencies for researchers should be interconnected. These facilities, if coordinated and centrally managed, have the capability to interconnect many computer networks into a single virtual computer network.
- As the second step, the existing computer networks that support research programs should be expanded and upgraded to serve 200-400 research institutions with 1.5 million bits per second capabilities.
- As the third step, network service should be provided to every research institution in the U.S., with transmission speeds of three billion bits per second.

A staged program of research and development can achieve the networking capability that is needed for the third step. This research and development effort will result in support to the U.S. research community and in an enhanced ability of the U.S. computer and communications industry to compete in world markets.

This report was conducted by an interagency group of the Committee on Computer Research and Applications of FCCSET. The report was requested by the 99th Congress in Public Law 99-383.

## 2. COMPUTER NETWORK STUDY

---

### 2.1. BACKGROUND

In 1986, the 99th Congress charged the Office of Science and Technology Policy (OSTP) with conducting "a study of critical problems and current and future options regarding communications networks for research computers, including supercomputers, at universities and federal research facilities in the United States" (Public Law 99-383, August 21, 1986). AT OSTP's direction, an interagency group under the auspices of the Federal Coordinating Council for Science, Engineering and Technology (FCCSET) was formed to carry out the computer network study. Agencies participating were DARPA, DoD, DOE, NASA, NBS, NSF, and NIH.

The Congress asked that the following issues be included in the study:

- the networking needs of the nation's academic and federal research computer programs, including supercomputer programs, over the next 15 years, including requirements in terms of volume of data, reliability of transmission, software compatibility, graphics capabilities, and transmission security;
- the benefits and opportunities that an improved computer network would offer for electronic mail, file transfer, and remote access and communications; and
- the networking options available for linking academic and research computers, including supercomputers, with a particular emphasis on the advantages and disadvantages of fiber optic systems.

This charge conveys the concerns of the Congress that effective network services for scientists may be approaching limits while requirements for access to networks are increasing. Computer networks are a vital support component for modern science, engineering and technology. Computer networks allow the large, diverse, and geographically dispersed U.S. research community to share large scale computing resources, to access remote research facilities, and to exchange information across great distances. Computer networks have the potential to support instantaneous communication and remote collaboration on a national and international scale. However, computer networks today cannot adequately support this communication and collaboration because of limited capacity and capability as well as lack of access to networks by all of the nation's academic, industrial and government research institutions.

In June 1985, the House Science and Technology Committee highlighted the importance of access to supercomputers by researchers at universities and laboratories. In 1985 FCCSET established a Network Working Group to coordinate Federal agency networking activities. A report "Interagency Networking for Research Programs" was published in February 1986 recommending the interconnection of existing Federally supported telecommunications networks and the formation of an Interagency Research Internet Organization.

In conducting its study during late 1986 and early 1987, the FCCSET Network Study Group enlisted the help of many experts from government, industry, and academia. White papers were invited on networking trends, requirements, concepts, applications, and plans. The group reviewed the status of existing research networks, analyzed the requirements of researchers to access networks, and assessed the capabilities of current technology. A workshop involving nearly 100 researchers, network users, network suppliers, and policy officials was held in February 1987 to air ideas, gather information, and

develop the foundation for the report to the Congress. The workshop participants discussed access requirements and future alternatives; special requirements for supercomputer networks; internetwork concepts; future standards and services requirements; security issues; and the role of government in networking.

## 2.2. FINDINGS

The information available to the Computer Network Study Group indicated that a strong, focused effort, supported by the Federal government, is needed to allow for adequate access to computing and to research facilities, to improve the state-of-the-art of computer networking, and to meet the challenge of foreign competition in this critical area of technology.

Today access to computer networks by researchers is haphazard and dependent upon individual funding or location. There is a great redundancy in the links from various agencies to each campus. Much broader coverage and better facilities are needed throughout the nation. High performance computers are partially driving the need for improved networking capabilities. They are capable of generating data much faster than it can be communicated using today's networking technology. The development of improved networking facilities can be compared to the development of the interstate highway system. Just as the interstate highway system stimulated economic development throughout the nation, so can data communications highways stimulate U.S. research and provide equitable access to resources.

## 2.3. REQUIREMENTS

Many scientific research facilities in the U.S. consist of a single, large, and costly installation such as a synchrotron light source, a supercomputer, a wind tunnel, or a particle accelerator. These facilities provide the experimental apparatus for groups of scientific collaborators located throughout the country. The facilities cannot be duplicated in all institutions because of cost. Wide area networks are the primary mechanism for making such facilities available nationwide. Examples include government-supported wide area networks such as ARPANET, HEPnet, MFENET, MILNET, NASnet, NSFnet, BITNET, and SPAN, as well as commercial facilities such as Tymnet and AT&T leased lines.

Today's networking resources are not adequate to support the needs of future U.S. researchers. Existing network links throughout the research community are generally low data rate (i.e., at most 56 kbit/s) and fully utilized. Some of these networks are severely overloaded, resulting in significant performance degradation. Additionally, more ubiquitous access is needed by the university research community, especially at smaller institutions. By 1990, U.S. researchers will need access to wide area networks that are one thousand times more capable than those available today. This estimate was based on analysis of existing network utilization, use of a typical site, experience with current local area networks, and expected future user populations. (See Volume II, Networking Requirements and Future Alternatives.) Remote high resolution interactive workstations will be essential for using computer graphics techniques which enable researchers to visualize and simulate two and three dimensional structures. Molecular biology, space exploration, cartography, ship and airplane design, and energy research applications are some of the research areas that would benefit from increased speed of data transfers. Higher speeds are also needed to allow sharing of large data bases produced by distributed research enterprises and to keep pace with future high performance computers.

Longer-range estimates vary (see volume II, Networking Requirements and Future Alternatives, and Internet Concepts), but it is clear that by the year 2000 the nation's research community will be able to make effective use of a high-capacity national network with speeds measured in billions of bits per second.

Without improved networks, speed of data transmission will be a limiting factor in the ability of future researchers to carry out complex analyses. Digital circuits are widely available today, at a transmission speed of 56 kilobits per second (kbit/s). For highly complex analyses such as examining molecular structures, investigating flows of gases and liquids, and conducting structural analyses, such speeds are impediments to productive work. Presenting computer generated images that appear to move requires 30 frames per second; each frame represents about 10 million bits per second (Mbit/s) of information. This presentation thus requires a transmission speed of 300 million bits per second of information. To support thousands of scientists simultaneously (even using advanced compression technology) would require backbone speeds of 300 billion bits per second (Gbit/s). See Appendix D for an example of collaborative research for which high speed networking is essential.

Within the next five years, Integrated Services Digital Network (ISDN) switched and non-switched circuits ranging from 64 kbit/s to 1.5 Mbit/s will be available in the larger metropolitan areas of the U.S. However, even these services will fall short of the requirements for computer networks. For example, by 1988 over 50 campus area networks (CANs) will be operational with advanced capabilities (100 Mbit/s). Wide area networks operating at a much slower data transmission rate (56 kbit/s to 1.5 Mbit/s) cannot handle the expected high data volume. See Figure 1.

Increased data communications capacity will be needed to support the effective use of supercomputers and high capacity work stations. While many scientists will have direct access to these facilities, networking will still be important for collaborative research that utilizes large programs and databases.

Other future requirements relate to interoperability and security. An individual scientist may find it necessary to interact with other scientists or machines on more than one network. Some of the networks are not compatible because they were developed according to design goals that did not include consideration of uses and technologies unrelated to the job at hand. Some of the networks are overloaded with traffic. Security is not uniformly good from network to network or from host to host.

## 2.4. RECOMMENDATIONS

The U.S. should undertake, as a national goal, the establishment of a National Research Network in a staged approach that supports the upgrade of current facilities, and development of needed new capabilities. Achievement of this goal would foster and enhance the U.S. position of world leadership in computer networking.

As rapidly as feasible, the National Research Network should be designed, deployed, and maintained as an advanced computer network. This network should interconnect substantially every academic, industrial and government research establishment and unique scientific resource to encourage scientific collaboration unhindered by distance and to permit the sharing of unique research facilities and resources. Since security of the network is a vital concern, appropriate policies should be adapted to protect the information in the network from threats, vulnerabilities and risks, and to assure a uniform level of security.

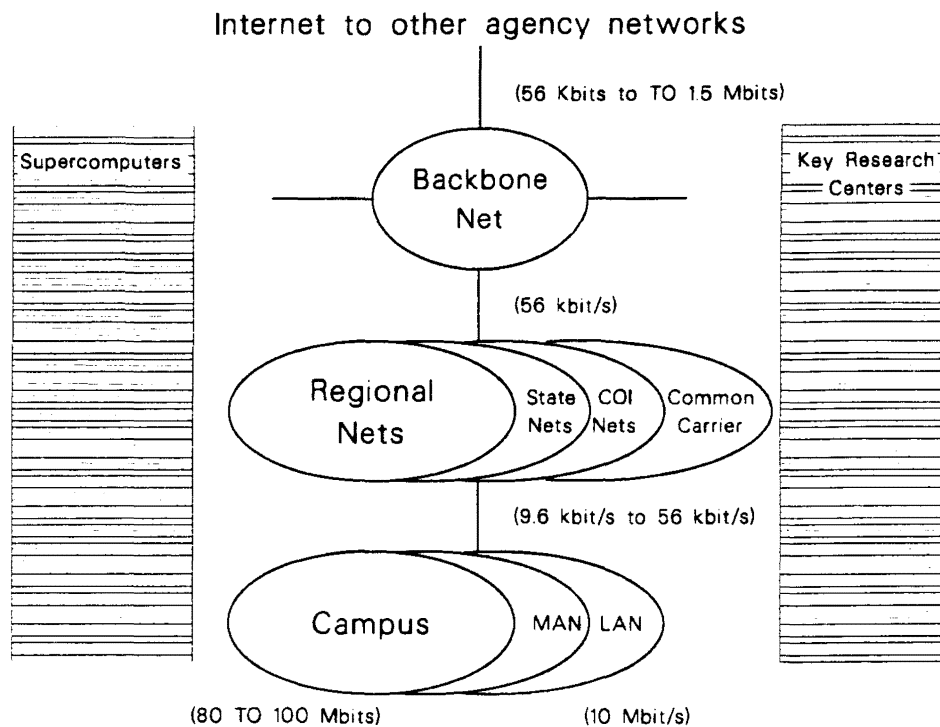


Figure 1. Research Network Hierarchy -- Today

Until the National Research Network can replace the current system, existing networks should be maintained and modified as they join the national network. Since supercomputer systems comprise a special and valuable national research resource with very high performance requirements, the responsibility for network access to supercomputers should be vested in the supercomputer centers themselves until the advanced computer network, capable of offering the requisite service level, is operational.

Industry should be encouraged through special incentives to participate in research, development, and deployment of the National Research Network. Tariff schedules which have been set for voice transmission should be re-examined in light of the requirements for transmission of data through computer networking.

To meet the goal for the National Research Network and to set an agenda for the future, the following actions are recommended:

- The Subcommittee on Computer Networking, Infrastructure and Digital Communications which was established by the Office of Science and Technology Policy on May 15, 1987, should oversee the first stage in the development and operation of the National Research Network, a coordinated internetwork that would include the Federal agencies that operate research supporting networks.
- The FCCSET Subcommittee on Computer Networking, Infrastructure and Digital Communications should identify a lead agency which would be responsible for requesting funds for the National Research Network, and eventually for selecting a contractor to manage the Network. The manager would be responsible for implementing Stages 2 and 3 of the National Research Network.

- As a first stage in the development of the National Research Network, the current Internet system developed by DARPA and networks supported by agencies should be interconnected over the next two years. These facilities, if coordinated and centrally managed, have the capability to interconnect many computer networks into a single virtual computer network. The Federal government should encourage and assist research facilities and academic institutions to establish local and campus area networks to connect to the Internet system. The estimated cost for this proposed upgraded service is \$5 million per year and should be implemented through the shared resources of NSF, DOE, DARPA, NASA, and HHS.
- In the second stage, new funding for development should be requested at \$5 million per year over the next five years to upgrade and expand the nation's existing computer networks, which support research programs, to achieve data communications at 1.5 Mbit/s to 200-400 U.S. research institutions. It is estimated that these expanded and upgraded facilities will require additional annual funding of approximately \$50 million to operate.
- In the third stage, a vigorous and focused program of research and development for the National Research Network should be immediately established. A total of \$400 million is needed over ten years to advance the knowledge base and technology of computer network capabilities in order to achieve data communications and switching capabilities to support transmission of three billion bits per second within fifteen years. These capabilities are 100,000 times more capable than currently available and will be essential to foster scientific collaboration and sharing of research resources. When fully deployed, the cost of operating this advanced network is estimated to be \$400 million per year, given the current commercial tariffs for data communications.

Support should be given to the development of standards and their harmonization in the international arena. Aggressive action is needed to increase user participation in the standards development process, to get requirements for standards expressed early in the development process, and to speed the implementation of standards in commercial off-the-shelf products. It is essential that standards development be carried out within the framework of overall systems requirements to achieve interoperability, common user interfaces to systems, and enhanced security.

## 2.5. BENEFITS

Implementation of the recommendations would address the issues that have been identified and would provide the U.S. scientific research community with a significant competitive advantage. Modernization of the nation's wide area networks by increasing speed, functionality, and size increases opportunities for research advances significantly. Greater network speed can reduce the time required to perform a given experiment and increase both the volume of data and the amount of detail that can be seen by researchers. Scientists accessing supercomputers would benefit particularly, because access speed is often critical in their work. Improved functionality frees scientists to concentrate directly on their experimental results rather than on operational details of the network. Increased network size extends these opportunities to tens of thousands of individuals located at smaller academic institutions throughout the nation. These modernization measures would significantly enhance the nation's competitive edge in scientific research.

The U.S. communications industries would also gain a significant competitive advantage. Development of modern, low-cost distributed computing facilities for wide area networks would help maintain the United States position of world leadership in utilization of wide area, high bandwidth networks. This would increase the nation's competitive edge in communications technology as well as scientific research. As a spinoff, it would help maintain the U.S. leadership position in computer architectures, microprocessors, data management, software engineering, and innovative networking facilities, and promote the development of international standards based on U.S. technology.

## 2.6. ACTION PLAN

The goal of the National Research Network interconnecting academic, industrial and government research organizations is reachable if we start now to support research and development on improved data communications speeds, to expand and upgrade existing networks, and to improve security and standards development.

This goal can be carried out in three stages, all of which must start immediately to achieve desired benefits. See Figure 2.

## 2.7. STAGE 1 THE INTERNET

This stage involves the internetworking and upgrade of existing agency networks. The various government networking activities touch a significant segment of the U.S. academic research community. The interlinking of some of these networks has already begun (e.g., NSFnet, the regional networks, the supercomputer networks, ARPANET, and other experimental defense networks). Most of these networks are adopting a common protocol suite to achieve interoperability. Through interagency collaboration, continued harmonization of protocols, and sharing of transmission facilities, these interlinked networks can be operational in two years. When these networks are in operation, performance will be 30 times that of today.

We recommend that each agency participating in the Internet (NSF, DOE, DARPA, NASA, and HHS) allocate \$1 million per year to accomplish the internetwork and that the FCCSET Subcommittee on Computer Networking, Infrastructure, and Digital Communications coordinate the activity.

The 1986 Report on Interagency Networking for Research Programs by the FCCSET Committee on Very High Performance Computing recommended the establishment of an organization under the direction of a FCCSET committee to provide overall coordination of the management and operation of an interagency network. The activities recommended in the report to carry out this stage of the development of the National Research Network are:

- establish, promulgate, and coordinate protocol standards and functional standards for the interagency internetwork;
- address issues of documentation and information availability between the involved agencies;
- coordinate interagency internetworking research projects.

About \$5 million per year spread over NSF, DOE, DARPA, NASA, and HHS is required to support this stage of development. Activities needed will be the purchase, installation, and operation of the major or 'core' network gateways between the existing and planned research networks; software development and maintenance, hardware maintenance, and operational monitoring and control of these gateways so that the interagency network is an available and reliable communications entity; installation of network routing, access control, and accounting procedures and tools, as these are developed; identification of the research and development projects necessary to create, maintain, and enhance the interagency network coordination of these projects with the constituent research network; implementation of standards.

## 2.8. STAGE 2 THE NATIONAL RESEARCH NETWORK

The goal of Stage 2 is to deliver network services of 1.5 Mbit/s to 200-300 research facilities. To provide this service, 45 Mbit/s speeds in optical fiber trunk lines must be achieved. This speed is needed to support computer graphics applications that enable users to visualize the results of calculations made on today's supercomputers and to provide the bulk capacity for thousands of users. This goal should be achievable through the application of sound development and engineering capabilities. About \$5 million per year is required for development of this phase of the National Research Network and about \$50 million per year to operate. A partnership with industry in the development of the National Research Network should be developed.

Private sector companies are offering an ever increasing array of communication services via satellite, recently installed optical fibers, microwave, and reorganized local service. Full advantage should be taken of these offerings as they change from time to time.

Fiber optic systems are most promising and are projected to operate at bandwidths which meet most of the requirements as defined by the U.S. research community. They also offer an additional advantage that, once installed, they should be able to accommodate more advanced, higher speed transmission equipment as it becomes available. However, lack of fiber optic ubiquity over the next decade may hinder its effectiveness to the end user or in reaching to the 'last mile'. In addition, satellite and digital microwave systems offer some economic and technical advantages which should not be overlooked for many requirements. For instance, satellite broadcast functionality may prove very beneficial to scientific collaborations and satellite transmission services may be the most cost effective approach for reaching less populated locations.

The limit of the current technology is very likely not bandwidth or connectivity. Researchers in the field suggest that the limitations will come first in gateways, routers, and switches and then later in the protocols and architectures of the networks. These issues must be addressed through a vigorous development effort to improve packet switching and protocols for networking.

## 2.9. STAGE 3 THE ENHANCED NATIONAL RESEARCH NETWORK

The goal of this stage is to deliver network services of 1.5 Mbit/s to every research facility in the U.S., and 1-3 Gbit/s to selected sites.

The technology to achieve this will require development and laboratory testing of new communications hardware, computer interfaces, transmission and routing protocols, and software design. The radically new designs that must result will require extensive laboratory and prototype testing.

The outcome of this process should be a design for a new national research network linking researchers and national support facilities such as supercomputer centers and research institutions. The first phase of deployment would involve settling the network design. Deployment of the trunks would follow, allowing interfaces to individual university campuses and research institutions. A national network to support research must be woven into the fabric of the national research infrastructure, and is as important as connecting major national research centers and facilities.

The estimated cost for research and development for this advanced facility is \$400 million over a ten year period, and about \$400 million per year may be needed to operate such a network. The cost of data communications will be a significant factor. Tariff structures created for voice communication are being imposed on data communication. The tariff structure should be reconsidered in light of the lower costs of high speed data communications using modern equipment.



The participation of industry in developing this network will be sought through the FCCSET Subcommittee on Networking, Infrastructure, and Digital Communications, the responsible agencies, and the contractor selected to operate the network. The participation of communication suppliers should be encouraged to provide low-cost fiber circuits during the critical ten-year research and development phase.




Figure 2. Timetable for a National Research Network.

## APPENDIX A - PUBLIC LAW 99-383

---

100 STAT.816

PUBLIC LAW 99-383—AUG. 21, 1986

### COMPUTER NETWORK STUDY

Research and  
development.  
42 USC 6614  
note.

SEC. 10. (a) The Office of Science and Technology Policy (hereinafter referred to as the "Office") shall undertake a study of critical problems and current and future options regarding communications networks for research computers, including supercomputers, at universities and Federal research facilities in the United States. The study shall include an analysis of—

(1) the networking needs of the Nation's academic and Federal research computer programs, including supercomputer programs, over the period which is fifteen years after the date of enactment of this Act, including requirements in terms of volume of data, reliability of transmission, software compatibility, graphics capability, and transmission security;

(2) 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; and

(3) the networking options available for linking academic and other federally supported research computers, including supercomputers, with a particular emphasis on the advantages and disadvantages, if any, of fiber optic systems.

Reports

(b) The Office shall submit to the Congress—

(1) within one year after the date of enactment of this Act, a report on findings from the study undertaken pursuant to subsection (a) with respect to needs and options regarding communications networks for university and Federal research supercomputers within the United States; and

(2) within two years after the date of enactment of this Act, a report on findings from the study undertaken pursuant to subsection (a) with respect to needs and options regarding communications networks for all research computers at universities and Federal research facilities in the United States.

## APPENDIX B - FCCSET COMMITTEE MEMORANDUM

---

THE WHITE HOUSE  
WASHINGTON

MEMORANDUM FOR THE FEDERAL COORDINATING COUNCIL ON  
SCIENCE ENGINEERING AND TECHNOLOGY (FCCSET)

SUBJECT: COMMITTEE ON COMPUTER RESEARCH AND APPLICATIONS

In FY 1986 the FCCSET Committee on High Performance Computing was chaired by Jim Decker, Office of Energy Research at DOE. The annual report is attached for your review. I would be pleased to hear any comments you may wish to make on this report.

In 1986, the congress charged OSTP (Public Law 99-383) to report by August 20, 1987 on critical problems and future options related to computer networks to support research in the United States. In order to carry out this study without disrupting the work of the existing committee I have revised the charter of the computer committee and renamed it to reflect a broader scope. A copy of that charter is attached. Jim Decker will chair the subcommittee on Scientific and Engineering Computing, Saul Amarel of DARPA will chair the subcommittee on Computer Research and Development, and Gordon Bell of NSF will chair the committee on Computer Networking, Infrastructure, and Digital Communications. The OSTP representative (currently Paul Huray) will coordinate the activities of the subcommittees and act as chairman of the Executive (Steering) Committee.

Sincerely,

William R. Graham  
Science Advisor to the President

# CHARTER OF THE COMMITTEE ON COMPUTER RESEARCH & APPLICATIONS OF THE FEDERAL COORDINATING COUNCIL FOR SCIENCE, ENGINEERING AND TECHNOLOGY

## BACKGROUND

The nation's need to maintain a strong national defense capability and to compete effectively in world trade is directly related to the advancement and application of computing and digital communications technology. Therefore, the U.S. Government must maintain technological leadership in computing and communications.

## PURPOSE

The Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) Committee on Computer Research and Applications shall address issues that relate to the retention of U.S. leadership in computing and digital communications particularly where government research, development, and procurement policies affect the advanced, high performance segments of the industry. It shall coordinate scientific and engineering applications and research in advanced computing and digital communications across the Federal Government. It shall maintain an awareness of government agency use of and research on advanced computing and communications in order to prevent undesirable duplication of effort and to share the benefits from the various agency initiatives. It shall monitor the infrastructure and manpower that support high performance computing and digital communications to ensure that the needs of the nation will be met.

The committee shall encourage and facilitate actions by government agencies to provide access to supercomputer facilities by researchers and to cooperate, where feasible, in interagency reciprocity in the sharing of advanced computer resources and communications networks. The committee will address and maintain an awareness of issues and technologies affecting networking between advanced facilities and users.

The committee shall encourage and facilitate actions by government agencies and government supported research performers to transfer newly acquired technology and expertise to the private sector as appropriate.

## IMPLEMENTATION

To achieve the stated purposes, three subcommittees shall be formed, dealing with:

- Scientific and Engineering computing.
- Computer Research and Development.
- Computer Networking, Infrastructure, and Digital Communications.

(Office automation and management information systems are not addressed by this charter). The subcommittees will have chairs appointed by the chair of the FCCSET committee.

## MEMBERSHIP

Membership of the subcommittees will include representation from the following agencies as appropriate:

Department of Commerce  
Department of Defense  
Department of Energy  
National Aeronautics and Space Administration  
National Science Foundation  
National Institutes of Health  
Intelligence Community

## LIAISON PARTICIPANTS

The subcommittee chairman may request participation of liaison members to serve as members of the subcommittees as they deem appropriate from:

National Academy of Science Computer Science  
and Technology Board  
Department of State  
Individual Services (Navy, Army, Air Force, DARPA)

## EXECUTIVE (STEERING) COMMITTEE

The chairmen of the three subcommittees, and a designated representative of the staff of OSTP (who shall be the committee chairman) shall form an executive committee to coordinate the activities of the subcommittees and to develop an appropriate plan of action.

## ADMINISTRATIVE PROVISIONS

- (a) The committee on Computer Research and Applications will report to the chair of the FCCSET through its chairman: its OSTP representative.
- (b) Meetings of the executive committee shall be called as deemed appropriate by members of the Executive Committee; the Director of OSTP or at the request of the FCCSET.
- (c) Special studies, analyses and recommendations may be initiated by the executive committee. As necessary, ad hoc subcommittees or working groups with participation not limited to the committee members may be formed to assist the committee in its work.
- (d) Member agencies will assign such working staff as requested by the subcommittee chairs and as is necessary and feasible for the conduct of committee activities. The respective agencies shall pay for the direct and incidental costs arising from the participation of their members and staff in committee activities.

## DURATION

The committee's activities and the continuing need for the committee shall be reviewed annually by the FCCSET.

## COMPENSATION

All members will be Federal employees who are allowed reimbursement for travel expenses by their agencies plus per diem for subsistence while serving away from their duty stations in accordance with standard government travel regulations.

## DETERMINATION

I hereby determine that the formation of the FCCSET Committee on Computer Research and Applications is in the public interest in connection with the performance of duties imposed on the Executive Branch by law, and that such duties can best be performed through the advice and counsel of such a group.

Approved:

May 15, 1987

\_\_\_\_\_  
(Date)

\_\_\_\_\_  
William R. Graham  
Science Advisor to the President  
Chair, Federal Coordinating  
Council for Science, Engineering  
and Technology

## FCCSET COMMITTEE ON COMPUTER RESEARCH AND APPLICATIONS

### Responsibilities of the Subcommittees

#### *Scientific and Engineering Computing*

Systems: Supercomputers to Workstations  
Graphics  
Performance: Benchmarks and Workloads  
Standards  
Applications  
Software and Algorithms  
Peripherals  
Supercomputing Access and Network Utilization  
Manpower

#### *Computer Research and Development*

Software Systems and Engineering  
Numeric and Symbolic Computing  
Algorithms and Theory  
Architecture  
AI and Robotics  
Database and Retrieval  
Manpower

#### *Computer Networking, Infrastructure, and Digital Communications*

Technologies and Research  
Systems  
Services  
Standards  
Interconnect and Coordination among National Networks  
Distributed Computing  
Manpower  
MOSIS Design and Manufacturing

## APPENDIX C - INDEX TO OTHER VOLUMES

---

This Computer Network Study is published in three volumes. Volume I contains the FCCSET recommendations on developing computer networks to support research in the U.S. Volume II contains the summaries of the February 1987 workshop discussions, which focused on six topics: access requirements and future alternatives, special requirements for supercomputer networks, internetwork concepts, future standards and services requirements, security issues, and the government role in networking. Volume III contains white papers that the Network Study Group invited on networking trends, requirements, concepts, applications, and plans.

The specific issues raised in Public Law 99-383 (August 21, 1986) are addressed in these three volumes. Following is an index to sections of Volumes I, II, and III that respond to the language of the Congressional charge:

### COMPUTER NETWORK STUDY

Research and development.  
42 USC 6614 note.

SEC 10. (a) The Office of Science and Technology Policy (hereinafter referred to as the "Office") shall undertake a study of critical problems and current and future options regarding communications networks for research computers, including supercomputers, at universities and Federal research facilities in the United States. The study shall include an analysis of—

(1) the networking needs of the Nation's academic and Federal research computer programs including supercomputer programs, over the period which is fifteen years after the date of enactment of this Act, including requirements in terms of volume of data, reliability of transmission, software compatibility, graphics capability, and transmission security;

### Requirements

#### Volume of data

Volume II,	Networking Requirements and Future Alternatives
Volume III,	Future Directions in Communications Research
	Local Area Networking with an Emphasis on Gateways and Digital Switches
	Networking Trends in Data Communications
	National Network Requirements
	Networking Requirements for Scientific Research
	Industry & Technology Trends
	DOE Networking Requirements

## Reliability of Transmission

- Volume II,                   Networking Requirements and Future Alternatives  
Computer Network Security
- Volume III,                 Future Standards and Service Requirements  
Future Directions in Communications Research  
Local Area Network Technology with  
Emphasis on Gateways and Digital Switches

## Software Compatibility

- Volume III,                 Advanced System Software for Supercomputers  
Impact of Distributed Functions on  
Network Requirements  
Network Requirements for Scientific Research

## Graphics Capabilities

- Volume I,                   Recommendations
- Volume III,                 The Role of the Graphics Workstation  
in Supercomputing  
National Networking Requirements

## Transmission Security

- Volume II,                 Computer Network Security
- Volume III,                 Future Directions in Communications Research

(2) 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; and

- Volume I,                   Recommendations
- Volume III,                 Networking, Some Observations on the Future  
The Role of Graphics Workstation in  
Supercomputing  
Network Requirements for Scientific Research



(3) the networking options available for linking academic and other federally supported research computers, including supercomputers with a particular emphasis on the advantages and disadvantages if any, of fiber optic systems.

- Volume I,                    Recommendations
- Volume II,                Internet Concepts  
                                  The Government Role in Networking
  
- Volume III,                Implementation Plan for Interagency Research Internet  
                                  The Federal Government's Role in  
                                  National Research Networking  
                                  The Role of the Government in  
                                  National Research Networks

Reports

- (b) The Office shall submit to the Congress—
  - (1) within one year after the date of enactment of this Act, a report on findings from the study undertaken pursuant to subsection (a) with respect to needs and options regarding communications networks for university and Federal research supercomputers within the United States; and

Volumes I, II, and III, FCCSET Report on Computer Networks to Support Research in the U.S. completed

- (2) within two years after the date of enactment of this Act, a report on findings from the study undertaken pursuant to subsection (a) with respect to needs and options regarding communications networks for all research computers at universities and Federal research facilities in the United States.

All issues have been addressed; no second year study is planned.

## APPENDIX D - NETWORKING REQUIREMENTS FOR MOLECULAR BIOLOGY

---

The science of molecular biology has made great strides in understanding and in manipulating fundamental life processes because of supercomputers and computer networking technology. Knowledge of molecular structure is of critical importance in the design of new drugs and treatment strategies. Currently, an intense effort is underway to deduce the molecular structure of the AIDS virus and, with this knowledge, there will exist a much better chance of developing strategies to combat it.

Most current work in this area relies on x-ray diffraction methods, which again relies on the ability to obtain the molecule under investigation in pure form and large amount. It is a much more difficult and computationally intensive task to approach this problem using the primary DNA sequence and then compute the composition and the shape of the protein molecule that the DNA codes represent (rather than measuring it by x-ray techniques). However, this computerized method is what must be done in order to cope with rising flood of DNA sequence data which is beginning to pour out of laboratories across the nation. In order to test and model the millions of possible molecular conformations in real time using graphical tools, computer cpu speeds of 200-500 Mflops will need to be complemented by networks capable of updating full screen, bit-mapped color images of molecules with real time performance which supports animation (refresh rates of 20-30 times per second). Network services 800-1000 Mbits/second to the end-user will be required for full implementation of such research systems by scientists on a nationwide basis.

To determine the three dimensional structures of biological macromolecules by x-ray crystallography, the x-ray diffraction data is expressed as tables of x-y-z coordinates of the component atoms the molecular structure. Using complex algorithms on supercomputer systems, atomic coordinates can be turned into three-dimensional color representations of these complicated molecules, many of which are made up of tens of thousands of atoms. Such color pictures of molecules can be manipulated, using joysticks or other pointing devices connected to scientific workstations, to reveal to the biologist areas of special biological importance. Using these techniques, it has been possible to determine by computer the probable sites for antibody formation against new and changing viruses, or the mutagenic effects of carcinogens on DNA.

The result of improved networks to the biological sciences will be an unparalleled new capacity to understand the estimated 100,000 cellular functions which govern the growth and development of human beings disposition to health and diseases. Molecular biology is leading us to a future where the computer is elevated from being an information provider to being a laboratory assistant, which is able to interpret questions, together with the available data, and to model the hypothesis being tested. But the staggering array and complexity of the molecules in living cells will pose an increasing requirement for supercomputer-based analysis methods, and for computer networks whose transmission speeds are several orders of magnitude higher than those which are now available to biological scientists in the U.S.

A Report to the Congress on  
Computer Networks to Support Research  
In the United States

\*\*\*\*\*

A Study of Critical Problems and Future Options

Volume II  
Reports from the Workshop on Computer Networks  
February 17-19, 1987, San Diego, California

Preliminary Issue

June 1987

Document prepared by the  
Computing and Communication Division,  
Los Alamos National Laboratory, Los Alamos, New Mexico.\*

This report was prepared as an account of work sponsored by agencies of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

\*The Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

FEDERAL COORDINATING COUNCIL ON SCIENCE, ENGINEERING AND TECHNOLOGY

---

A Report to the Congress on  
Computer Networks to Support Research  
In the United States

\*\*\*\*\*

A Study of Critical Problems and Future Options

Volume II  
Reports from the Workshop on Computer Networks  
February 17-19, 1987, San Diego, California



## Contents

---

<b>Introduction to the Study</b> .....	1
<b>Agenda--Workshop on Computer Networks</b> .....	3
<b>Workshop Participants</b> .....	4
<b>1. Networking Requirements and Future Alternatives</b> .....	7
1.1. Background .....	7
1.2. Approach .....	8
1.3. Network Demand Summary .....	8
1.4. Supply .....	10
1.5. Alternatives .....	11
1.5.1. Supplying Capacity .....	11
1.5.2. Improved Usability .....	12
1.6. Conclusions .....	12
1.7. Benefits .....	14
1.8. Working Group Participants .....	15
1.9. Principal Networking Sites .....	16
1.10. References .....	20
<b>2. Special Requirements for Supercomputer Networks</b> .....	21
2.1. Requirements .....	21
2.2. Status .....	23
2.2.1. Category I: Mission-Oriented Networks .....	23
2.2.2. Category II: Confused Orientation .....	23
2.2.3. Category III: Serving all Comers .....	24
2.2.4. Category IV: Ad Hoc Solutions .....	24
2.2.5. Summary of Category Use .....	24
2.3. Conclusions .....	24
2.4. Working Group Participants .....	25
<b>3. Internet Concepts</b> .....	27
3.1. Summary of Current Research Networking in the U.S. ....	27
3.1.1. DoD Networks .....	27
3.1.2. NSF Networks .....	28
3.1.2.1. CSNET Description .....	28
3.1.2.2. NSFnet Description .....	28
3.1.3. DOE Networks .....	29
3.1.3.1. MFENET/ESnet .....	29
3.1.3.2. HEPnet/LEP3NET .....	29
3.1.4. NASA Networks .....	31
3.2. The Academic Internet of 1992 .....	32
3.2.1. Major Conclusions .....	32
3.2.2. Technology Issues for the 1992 Research Internet .....	33
3.2.3. Opportunities for Resource Sharing and Coordination .....	34

3.2.4. Summary .....	34
3.3. Academic Networking in the Year 2000 .....	35
3.3.1. User Requirements .....	35
3.3.1.1. Overall Number of Users .....	35
3.3.1.2. Classes of Traffic .....	35
3.3.1.3. Required Bit Rate .....	36
3.3.1.4. Switching .....	36
3.3.1.5. Inhomogeneity .....	36
3.3.2. Resources/Technology Available .....	36
3.3.2.1. Bandwidth .....	37
3.3.2.2. Switching .....	37
3.3.2.3. Role of Satellites .....	38
3.3.2.4. Communication Software .....	38
3.3.2.5. Local Processing .....	38
3.3.3. Conclusions .....	38
3.3.4. Caveat .....	39
3.4. Summary .....	39
3.4.1. Considerations to the Year 1992 .....	39
3.4.2. Considerations to the Year 2000 .....	39
3.5. Working Group Participants .....	40
<b>4. Future Standards and Service Requirements .....</b>	<b>41</b>
4.1. The Importance of Standards .....	41
4.2. The Economics of Standards .....	41
4.3. History and Trends in Network Services and Standards .....	42
4.3.1. History of Computer Networking Standards .....	42
4.3.2. Trends in Standards Development .....	46
4.4. Lessons Learned in Networking Standardization .....	47
4.5. Networking Services for the Scientific Community .....	49
4.6. Standard Networking Services .....	49
4.7. Service Reliability .....	51
4.8. Transition Strategies .....	53
4.9. Working Group Participants .....	54
4.10. References .....	55
<b>5. Computer Network Security .....</b>	<b>57</b>
5.1. Network Security Objectives .....	57
5.2. General Observations .....	57
5.3. Threats, Vulnerabilities, and Risks .....	58
5.4. Research Computer Network Security Program .....	59
5.5. Research Topics .....	59
5.5.1. Personal Authentication .....	60
5.5.2. Access Control .....	60
5.5.3. Integrity Assurance .....	60
5.5.4. Confidentiality Assurance .....	60



5.5.5. Security Event Monitoring .....	61
5.6. Verification .....	61
5.7. Summary and Conclusions .....	61
5.8. Security Policy Examples: Research Networks .....	62
5.9. Working Group Participants .....	63
<b>6. The Government Role in Networking .....</b>	<b>65</b>
6.1. Introduction .....	65
6.2. Summary .....	65
6.3. Full Deployment of Current Technology Networks .....	66
6.4. Research to Develop the Next Generation Network .....	69
6.5. Deployment of the Next Generation Network .....	71
6.6. Working Group Participants .....	72
6.7. References .....	73
<b>Glossary .....</b>	<b>75</b>
<b>Index .....</b>	<b>83</b>



## INTRODUCTION TO THE STUDY

---

The 99th Congress charged the Office of Science and Technology Policy (OSTP) with conducting "a study of critical problems and current and future options regarding communications networks for research computers, including supercomputers, at universities and federal research facilities in the United States" (Public Law 99-383, August 21, 1986). At OSTP's direction, an interagency group under the auspices of the Federal Coordinating Council for Science, Engineering and Technology (FCCSET) was formed to carry out the study.

The Congress asked that the following issues be included in the study:

- the networking needs of the nation's academic and federal research computer programs, including supercomputer programs, over the next 15 years, including requirements in terms of volume of data, reliability of transmission, software compatibility, graphics capabilities, and transmission security;
- the benefits and opportunities that an improved computer network would offer for electronic mail, file transfer, and remote access and communications; and
- the networking options available for linking academic and research computers, including supercomputers, with a particular emphasis on the advantages and disadvantages of fiber optic systems.

Networks are essential components of science, engineering, and technology, making it possible for researchers to share resources, access remote facilities, exchange information and software, and improve their collaborative relationships. The challenge is to make effective use of existing computer networks and to create advanced capabilities for the future. This will not only enhance scientific research in many disciplines, but will also help to advance the state-of-the-art of U.S. networking technology.

In conducting the network study during late 1986 and early 1987, the FCCSET group enlisted the help of many experts from government, industry, and academia. White papers were invited on networking trends, requirements, concepts, applications, and plans. A workshop involving nearly 100 researchers, network users, network suppliers, and policy officials was held in February 1987 to air ideas, gather information, and develop the foundation for the report to the Congress. Industry organizations were invited to provide information on the expected costs and benefits of fiber optic systems for networks.

As a result of the collaborative efforts of many dedicated people, the report to the Congress has been completed. It is published in three volumes:

- Volume I contains the FCCSET recommendations on developing computer networks to support research in the U.S.
- Volume II contains summaries of the February 1987 workshop discussions, which focused on six topics: access requirements and future alternatives, special requirements for supercomputer networks, internetwork concepts, future standards and services requirements, security issues, and the government role in networking.
- Volume III contains the invited white papers.

The workshop summaries in Volume II and the white papers in Volume III are presented as developed by their authors. No attempt has been made to achieve unanimity of opinion; there are many points of view expressed on a variety of network-related subjects.

I gratefully acknowledge the participation and support of the many people who have contributed to this report--the workshop participants; the chairs of the workshop groups; the San Diego Supercomputer Center, which hosted the workshop; the authors of the white papers; the staff of the Los Alamos National Laboratory who prepared the three volumes of the *Report to the Congress on Computer Networks to Support Research in the U.S.*; and the members of the FCCSET group that conducted the study.

My special thanks to John Cavallini for his support to me in organizing this study and preparing the report.

James H. Burrows  
Chair  
Computer Network Study

Computer Network Study Group

\*\*\*\*\*

Ron Bailey, National Aeronautics and Space Administration  
Gordon Bell, National Science Foundation  
James Burrows, National Bureau of Standards  
John Cavallini, Department of Energy  
Michael Corrigan, Department of Defense  
Paul Huray, Office of Science and Technology Policy  
Thomas Kitchens, Department of Energy  
James Oberthaler, Department of Health and Human Services  
Dennis Perry, Defense Research Projects Agency  
Shirley Radack, National Bureau of Standards  
Rudi Saenger, Naval Research Laboratory  
Daniel Van Belleghem, National Science Foundation  
Stephen Wolff, National Science Foundation

## AGENDA--WORKSHOP ON COMPUTER NETWORKS

---

Holiday Inn Embarcadero  
San Diego, California  
February 17-19, 1987

### Monday, February 16

3:00 p.m. - 7:00 p.m.           Registration, Lobby Foyer

### Tuesday, February 17

9:00 a.m. - 10:00 a.m.           Introduction to Workshop - James Burrows and Gordon Bell

10:00 a.m. - 12:00 p.m.           Working Group Meetings in rooms to be assigned

Group A - Internet Concepts Chair: Lawrence Landweber

Group B - Networking Requirements and Future Alternatives Chair: Sandy Merola

Group C - Future Standards and Services Requirements Chair: Richard desJardins

Group D - Security Issues Chair: Dennis Branstad

Group E - Government Role in Networking Chair: Jesse Poore

Group F - Special Requirements for Supercomputer Networks  
Chair: Robert Borchers

12:00 p.m. - 1:00 p.m.           Luncheon buffet at Holiday Inn

1:00 p.m. - 5:00 p.m.           Continuation of Group Meetings

### Wednesday, February 18

8:30 a.m. - 12:00 p.m.           Group Meetings

12:00 p.m. - 1:00 p.m.           Luncheon buffet at Holiday Inn

1:00 p.m. - 5:00 p.m.           Continuation of Group Meetings - Development of outline, summaries,  
and recommendations by each group

5:30 p.m. - 7:00 p.m.           Tour of San Diego Supercomputer Center; bus transportation available.

### Thursday, February 19

8:00 a.m. - 12:30 p.m.           Working Group summary presentations

12:30 p.m. - 1:30 p.m.           Luncheon at Holiday Inn

1:30 p.m. - 5:00 p.m.           Discussion of Working Group reports and development of final report  
by Planning Group

## WORKSHOP PARTICIPANTS

---

F. Ron Bailey, NASA Ames Research Center  
C. Gordon Bell, National Science Foundation  
Robert R. Borchers, Lawrence Livermore National Laboratory  
Bob Braden, University of Southern California - ISI  
Dennis Branstad, National Bureau of Standards  
Alison Brown, Cornell University  
James H. Burrows, National Bureau of Standards  
Roger M. Callahan, National Security Agency  
John S. Cavallini, Department of Energy  
Vinton Cerf, NRI  
Roger Cheung, Hewlett Packard  
Charles Crum, National Cancer Institute  
Henry Dardy, Naval Research Laboratory  
John Day, Codex/Motorola  
Richard desJardins, Computer Technology Associates  
Alan Demmerle, National Institutes of Health  
Debra Deutsch, Bolt Beranek and Newman Inc. Laboratories  
P.W. Dillingham, Jr., Cray Research, Inc.  
Dan Drobnis, San Diego Supercomputer Center  
Dennis Duke, SCRI-Florida State University  
Richard Edmiston, Bolt Beranek and Newman Inc. Laboratories  
David Farber, University of Delaware  
A. Frederick Fath, NASA/Boeing  
John Fitzgerald, LLNL/National Magnetic Fusion Energy Center  
Dieter Fuss, Lawrence Livermore National Laboratory  
David Golber, UNISYS  
David A. Gomberg, Mitre Corp.  
James L. Green, NASA  
Paul E. Green, IBM Research  
Phill Gross, Mitre Corp.  
Dennis Hall, Lawrence Berkeley Laboratory  
Ronald J. Ham, Digital Equipment Corp.  
Jack Haverly, Bolt Beranek and Newman Inc. Communications  
Paul G. Huray, Office of Science and Technology Policy  
Bill Johnson, Digital Equipment Corp.  
William Prichard Jones, NASA Ames Research Center  
Robert E. Kahn, National Research Initiatives  
Sid Karin, San Diego Supercomputer Center  
Jonathan S. Katz, The Analytic Sciences Corp.  
Charles M. Kennedy, U.S. Army Ballistic Research Laboratory  
Steve Kent, Bolt Beranek and Newman Inc. Communications  
John Killeen, LLNL/National Magnetic Fusion Energy Center  
Tom Kitchens, Department of Energy  
K. L. Kliever, Purdue University  
Larry Landweber, University of Wisconsin  
Thomas Lasinski, NASA Ames Research Center  
Tony Lauck, Digital Equipment Corp.  
Larry Lee, Cornell National Supercomputer Center  
James Leighton, LLNL/National Magnetic Fusion Energy Center  
Barry Leiner, NASA Ames Research Center  
Michael Levine, Pittsburgh Supercomputing Center

Dan Lynch, Advanced Computing Environment  
Richard Mandelbaum, University of Rochester  
Daniel Masys, M.D., National Institute of Health  
Noel Matchett, Information Security Inc.  
Rivi Mazumdar, Columbia University  
Fred McClain, San Diego Supercomputer Center  
Patrick McGregor, Contel Business Networks  
Sandy Merola, Lawrence Berkeley Laboratory  
Hugh Montgomery, Fermi National Accelerator Laboratory  
John Morrison, Los Alamos National Laboratory  
Norman Morse, Los Alamos National Laboratory  
Sushil Munshi, United Telecom. Inc.  
Dan Nessett, Lawrence Livermore Laboratory  
Harvey Newman, CalTech  
Jim Oberthaler, Department of Health and Human Services  
Ari Ollikainen, NASA Ames Research Center  
Dennis G. Perry, Defense Advanced Research Projects Agency/ISTO  
Abraham Peled, IBM Watson Research Center  
Richard Pietravalle, Digital Equipment Corp.  
Jesse H. Poore, University of Tennessee  
Jon Postel, University of Southern California/ISI  
Douglas R. Price, Sparta. Inc.  
Shirley M. Radack, National Bureau of Standards  
Harry L. Reed, Jr., Ballistic Research Laboratory  
Glenn Ricart, University of Maryland  
Richard T. Roca, AT&T Bell Labs  
Paul Rotar, National Science Foundation  
Karen Roubicek, Bolt Beranek and Newman Inc. Laboratories  
Stan Ruttenberg, University Corporation for Atmospheric Research  
Rudi F. Saenger, Naval Research Laboratory  
Marvin Sirbu, Carnegie Mellon University  
Miles Smid, National Bureau of Standards  
Ted Spitzmiller, Los Alamos National Laboratory  
David Staudt, National Science Foundation  
Dave F. Stevens, Lawrence Berkeley Laboratory  
George Sullivan, Defense Communications Engineering Center  
Joseph Tardo, Digital Equipment Corp.  
Joseph Timko, AT&T Bell Labs  
Dan Van Belleghem, National Science Foundation  
Stephen Walker, Trusted Information Systems  
Al Weis, IBM Research  
Bob Wilhelmson, National Center of Super Applications  
James Wilson, House Committee on Science, Space, and Technology  
Kenneth G. Wilson, Cornell University  
Stephen Wolff, National Science Foundation  
John G. Zornig, Digital Equipment Corp.





# 1. NETWORKING REQUIREMENTS AND FUTURE ALTERNATIVES

---

Sandy Merola  
Lawrence Berkeley Laboratory

## Abstract

*The Working Group on Networking Requirements and Future Alternatives recommends creation of an international, interagency networking facility for science, whose 15-year mission is*

- *to ensure that U.S. scientists have available the most advanced wide area networking facilities in the world, and*
- *to ensure that U.S. wide area network technology maintains a position of world leadership.*

*A minimum of 1.5-Mbit/s access to major government and academic research centers should be provided. Such a network would greatly benefit the competitive position of the United States in scientific research. It would also place the U.S. in a leadership position in utilization of high bandwidth, wide area networks. United States industries supporting wide area network technologies would gain a significant competitive advantage over other countries. An ongoing program of research and development into both wide area network technology and network management is necessary for this endeavor to be successful.*

*As part of the second year study, the Working Group recommends that an interagency coordinating committee be established to identify short-term implementation issues that can be investigated and resolved in parallel with long-term issues. This would provide immediate benefit to the nation's scientific community.*

## 1.1. BACKGROUND

Many scientific research facilities in the U.S. consist of a single, large, and costly installation such as a synchrotron light source, a supercomputer, a wind tunnel, or a particle accelerator. These facilities provide the experimental apparatus for groups of scientific collaborators located throughout the country. The facilities cannot be duplicated in all states because of cost. Wide area networks are the primary mechanism for making such facilities available nationwide. Examples include government-supported wide area networks such as ARPANET, HEPnet, MFENET, MILNET, NASnet, NSFnet, SPAN, and so on, as well as commercial facilities such as Tymnet, BITNET, and AT&T leased lines. The cost of such networks is generally much less than the cost of the research facility.

Congress recently enacted legislation calling for an investigation of the future networking needs over the next 15 years for the nation's academic and federal research computer programs. The Federal Coordinating Council on Science, Engineering and Technology (FCCSET) formed a Network Study Group to coordinate investigation of the benefits, opportunities for improvements, and available options with particular attention to supercomputing. Within the Network Study Group, the Working Group on Network Requirements and Future Alternatives was formed to identify network demand during the next 5 years and to recommend a strategy for meeting that demand. This document is the Working Group's report.

## 1.2. APPROACH

The following approach was taken.

- The networking plans of the U.S. research community were analyzed, so that a 5-year network demand summary can be created.
- Corporations that provide telecommunications services were surveyed, with particular attention to the possible use of fiber optics and related cost/capacity gains.
- Issues related to interagency sharing of network facilities were identified.
- Alternative methodologies for meeting total network demand were considered.
- A 5-year networking strategy was developed and presented to the FCCSET Network Study Group.

## 1.3. NETWORK DEMAND SUMMARY

Four methods of estimating network demand were used.

- **Analysis of existing network utilization:** Wide area networks are used by scientists to access unique remote facilities (supercomputers, accelerators, analysis software, and databases) and as a critical mechanism for communication and coordination among the large geographically distributed U.S. and international scientific collaborations (Figure 1 and Section 1.9). High-speed local area networks are being connected to low-speed wide area networks throughout the research community. Communication speeds of 1.5 Mbit/s, digital data service (DDS), and packet networks have been introduced to wide area networks, and their use has become widespread. Nevertheless, wide area networking capacity has not kept up with the capacity of local area networks. Some wide area networks handle both high data volume and highly interactive traffic over the same communications links. This results in suboptimal performance. At the functional level, wide area network user interfaces have not kept up with their counterparts in local area networks.

The Working Group heard presentations of current and planned networking in the Department of Defense, Department of Energy, National Aeronautics and Space Administration, and the National Science Foundation (NSF). Many scientific research centers funded by these agencies are physically connected to more than one network. The backbones for the major networks are similar in topology, and existing network links throughout the community are generally fully utilized. Some of these networks are severely overloaded, resulting in significant performance degradation. Additionally, more ubiquitous access is needed by the university research community, especially at smaller institutions. For example, there is a clear unmet need for nationwide, high-speed access to large scientific databases. The Working Group noted that in many cases demand for capacity seriously exceeded current supply.<sup>1-3</sup>

- **Estimation based on typical site:** A direct estimation of network demand was made using a major NSF university site as a basis. Network usage included wide area network facilities for supercomputer access as well as an extensive local area network. An absolute level of network demand for the next 5 years was estimated using three different models: task, user, and external flow. The task model focused on the network load generated by typical network tasks. The user model identified demand as a function of typical university network users. The external flow model centered on the university as an entity and estimated networking demand between it and other external locations. The three values of predicted network traffic were in agreement within an order of magnitude. They indicated a thousandfold increase in needed capacity over current network resources.<sup>4</sup>

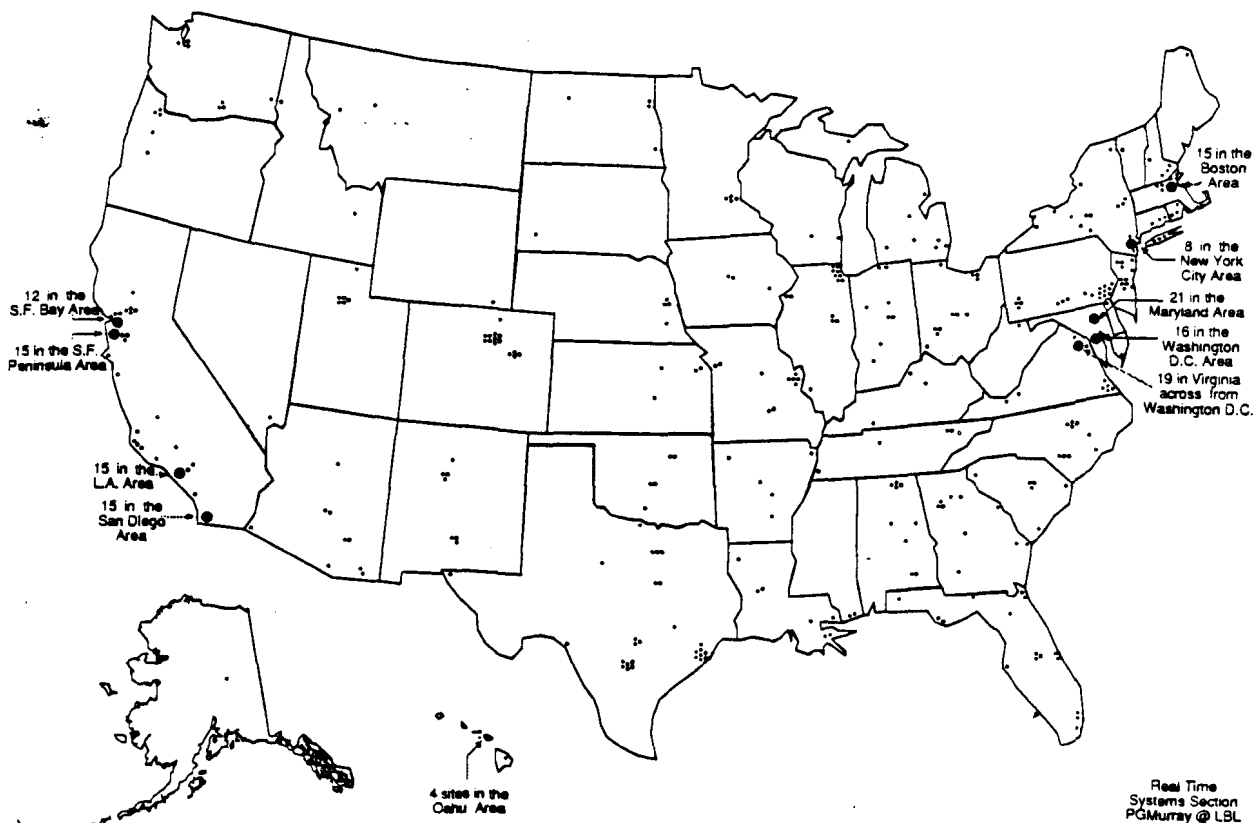


Figure 1. Principal networking sites - see Section 1.9 for a listing

- **Extrapolation from experience with local area networks:** This method also projected need for a thousandfold increase in wide area network capacity over the next 5 years. A remote supercomputer access scenario was presented to demonstrate how network transparency can increase the speed and accuracy with which engineering decisions can be made. It was argued that one order of magnitude is needed to create a nationwide distributed file system on an existing 56-kbit/s network; another order of magnitude is needed to provide interactive monochrome graphics<sup>5, 6</sup> and a third order of magnitude is needed to accommodate expected increases in basic computer speeds. As more users are added, further increases in demand are anticipated.
- **Estimation based on expanded user community:** The above analyses estimate load increases for existing network topologies. There is an important additional need to extend network service to the smaller universities throughout the nation. This would add another factor of 2 to 3 to the above estimates. Since by definition these research sites are not currently connected to an existing wide area network, this represents a demand for more communications lines rather than an increase in line speeds.<sup>1</sup>

There is a further need to extend network service to international sites. Access to overseas scientific collaborations would significantly enhance the quality of U.S. science by providing researchers with access to remote experimental apparatus, data, and personnel. It would also enhance U.S. prestige in the scientific research community by providing overseas collaborators with access to U.S. facilities, data, and personnel. The effect on network traffic would be negligible, but network size would be increased dramatically.

### 1.4. SUPPLY

Several major U.S. telecommunications corporations were represented on the panel. They jointly provided a summary of expected industry-wide technological trends over the next 5 years.<sup>7-10</sup> Cost/capacity forecasts and opportunities for use of fiber optic technology in the U.S. scientific research community were also presented.

The leading trends in U.S. telecommunications technology are the decreasing cost of component materials and the widespread, though not ubiquitous, availability of fiber optics (Figure 2). The transport capabilities of the U.S. telecommunications industry will greatly increase during the next 5 years, as witnessed by the following observations. Packet switching rates are expected to rise to 10,000 packets per second (25 Mbit/s). Digital circuits are widely available at 56 kbit/s today. Within the next 5 years, Integrated Services Digital Network (ISDN) switched and nonswitched circuits ranging from 64 kbit/s to 1.5 Mbit/s will be available in the larger metropolitan areas of the U.S. The digital interexchange transmission rates available to users are at 1.5 Mbit/s in general and will rise to 45 Mbit/s between larger metropolitan areas. Services of 150 Mbit/s could be made available by special arrangement. ISDN 64-kbit/s service will be present in about 20% of the U.S. market by the end of the 5-year period. The ability of the user to customize service (such as time of day conversion and simultaneous coordinated voice and data), as well as the availability and general use of applications services (such as X.400 mail and electronic document interchange) will dramatically increase.

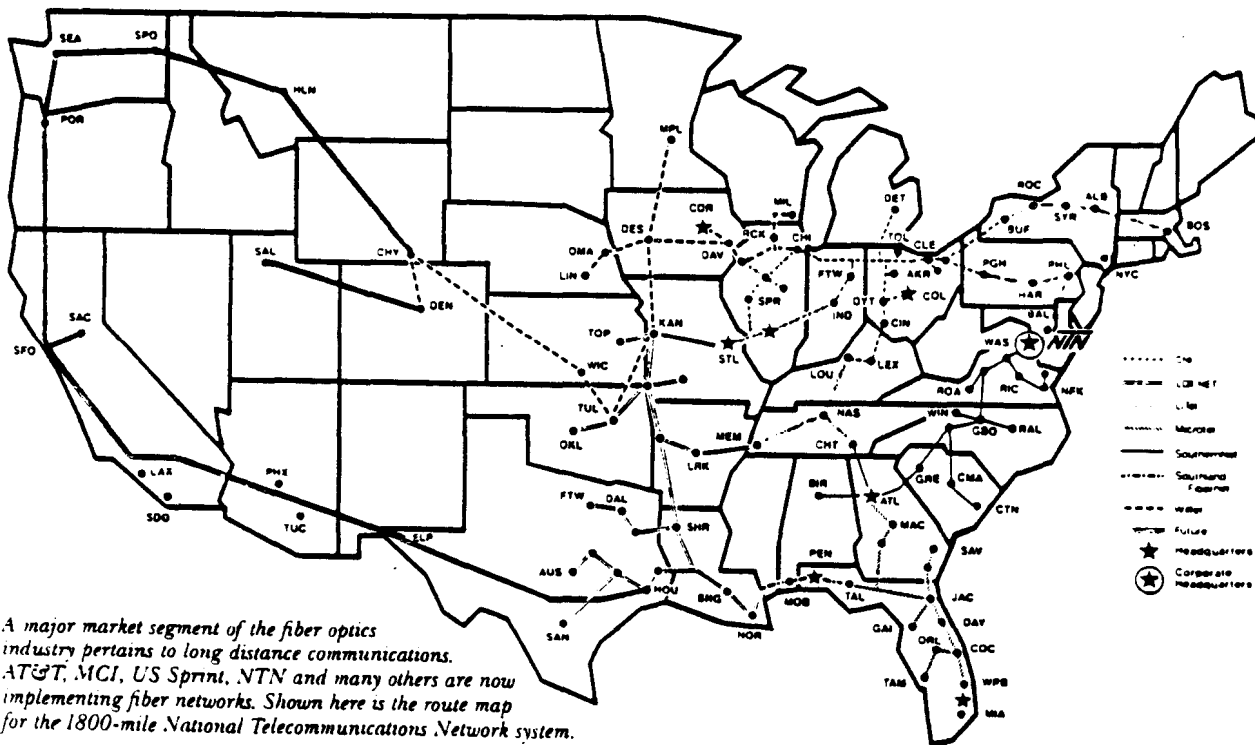


Figure 2. Fiber optic network links.<sup>11</sup>

Fiber optic technology is driving media costs downward. The cost of basic private line telecommunications services could fall by a factor of 20% to 50% during the upcoming 5 years. Any expectation that fiber would more dramatically reduce costs to the typical telecommunications user must be balanced by the recognition that the fiber itself is only one component of total transmission service cost.

It was recognized that the combination of fiber optic technology and the large amount of aggregate interagency demand may offer the scientific research community unique opportunities to acquire increasingly cost-effective bandwidth. This is only possible in the case of a long-term lease of very high bandwidth circuits. This ensures industry recovery of capital investment costs. If such a national network infrastructure were established as a long-term interagency goal, migration to such a topology is possible using existing standard telecommunications technologies, including satellite, microwave, copper, and fiber optic transmission media.

## 1.5. ALTERNATIVES

### 1.5.1. Supplying Capacity

The need to increase wide area network capacity by a thousandfold is justified both by increased opportunities for scientific breakthroughs and by the need to maintain the nation's position of world leadership in wide area network technology. Although industry projections indicate the necessary bandwidth will certainly be available as a national backbone, the required bandwidth will not be available all the way to the end user's site. The Working Group felt the most cost-effective way to proceed would be to provide the needed bandwidth in stages.

The Working Group recognized that a factor of 30 improvement could be achieved simply and cost effectively by:

- (1) tuning existing protocol implementations and managing access,
- (2) installing smarter congestion control algorithms,
- (3) upgrading existing 56-bit/s trunks to 1.5- and 45-Mbit/s lines in a judicious manner, and
- (4) providing type-of-service routing for efficient performance on high data volumes as well as highly interactive traffic.<sup>3</sup>

Beyond that, another factor of 30 is needed to meet the projected demand. The Working Group identified two promising approaches:

- (1) develop more optimal distribution of network services between user systems and server systems to make more efficient use of the available bandwidth, and
- (2) develop powerful gateway computers that compress data entering wide area networks and decompress the data at its destination. Such machines could also provide encryption without significant additional overhead.

The two approaches are entirely complementary. Thus, each might contribute a factor of 5 or 6, for a combined factor of 30. However, optimal distribution software is not available today, and data compression computers are only available for video compression. Therefore, applied research in these and other promising approaches is required.

### 1.5.2. Improved Usability

The Working Group agreed that an interagency, international network would significantly enhance the U.S. scientific research environment. To ensure ease of use, some peripheral issues must be addressed.

- **Global management and planning:** The ARPANET provides valuable experience in operating connected networks without global management. For example, ARPANET management reported that traffic generated by external networks created internal performance problems that are unmanageable. Similarly, inefficient protocol implementations cannot be prevented, since no central authority exists. This results in reduced network performance for all users. ARPANET management concluded that global management is essential to provide guaranteed performance. The Working Group agreed with this conclusion.
- **User services:** Consulting help and documentation are necessary for any facility accessed directly by end users. However, most scientists are not interested in networks per se, but only in the resources they make available. If a network could be made transparent or nearly so, the need for consulting help and documentation would be significantly reduced.
- **Reliability:** A wide area network in scientific research must be more reliable than many existing networks because of its critical role in supporting operation of remote experiments.
- **Extensibility:** The network will grow significantly in the next 15 years. It must be possible to expand it incrementally and to join it with other networks, both national and international.
- **Evolutionary:** To prevent obsolescence, the network must be tolerant of change. It must be designed in such a way that new protocols and services can be added without significantly disrupting existing services. This ensures the nation's scientists will keep a competitive edge in advanced networking technology. The rich environment for development of new products ensures that the technology itself maintains a competitive edge.

## 1.6. CONCLUSIONS

Five major conclusions about future networking requirements were drawn by the working group.

- (1) An interagency scientific network facility should be created whose 15-year mission is
  - to ensure that U.S. scientists have available the most advanced networking facilities in the world, and
  - to ensure that U.S. wide area network technology maintains a position of world leadership.
- (2) A phased implementation plan should be developed to provide these advanced network facilities to the nation's scientists. Rough guidelines should be to increase the effective capacity of existing networks tenfold in 3 years, a hundredfold in 5 years, and a thousandfold in 10 years.
  - Existing wide area scientific networks should be overhauled to provide 56-kbit/s service to end users at about 30% of maximum load. Trunk lines of 1.5 to 45 Mbit/s would be necessary in some areas to provide the needed bandwidth to end users. Existing protocol implementations should be checked and tuned to eliminate unnecessary congestion from inefficient implementations. Networks from all U.S. government agencies funding academic and federal scientific research would be upgraded.
  - Modern networking facilities such as wide area network file systems, distributed scientific databases, distributed window systems, and distributed operating systems should be developed

and installed, along with facilities for users to find and use network resources from remote sites. Existing communications facilities should be upgraded tenfold to 1.5 Mbit/s to end users as necessary to handle anticipated increases in load. Very high bandwidth trunk lines may be necessary in some areas to provide the needed 1.5-Mbit/s service to end users.

- More advanced facilities such as wide area color graphics capabilities and remote control of experiments should be developed and introduced. Existing communications capacity should be upgraded tenfold to handle the load increase by using hardware and software technology developed as a result of applied research.
  - To handle an anticipated increase in hardware speeds, existing communications links should be upgraded another tenfold as newer and faster computers become available in the mid 1990s.
  - New local area network facilities should be tracked so that the more promising new products can be made available in wide area networks.
  - Coverage should be expanded so that most colleges and universities in the U.S. will have access to the network in 5 years, with the remainder having access in 10 years.
- (3) An applied research and development program in advanced communications and network techniques should be implemented to provide the following.
- Provide the technology needed to increase the effective bandwidth of communications links would involve
    - more optimal distribution of functions between local hosts and remote hosts to minimize the need for raw network bandwidth,
    - high-performance systems that compress data entering a wide area network and decompress it at its destination,
    - development of gateway technology in general, and
    - utilization of formal language theory and other innovative techniques to design components that fail in a diagnosable manner.
  - Provide better ways to access remote resources that are needed to increase opportunities for scientific breakthroughs. Local area networks are the only cost-effective testbed for such facilities today. As capacity of wide area networks increases, a new source for network innovations can be expected to emerge.
  - Provide better tools and techniques for management of networks as needed.
- (4) An ongoing basic research program into future network architectures to ensure continued leadership in use of scientific networks, as well as national leadership in wide area network technology, is necessary.
- (5) The panel recommends that issues of network design, cost analysis, management authority, and implementation be addressed by the second year study. Within this framework, an interagency coordinating committee should be established to identify issues that can be investigated and resolved in the short term. An important short-term issue is implementation of the first factor of 30 improvement to existing networks. This can provide immediate benefit to the nation's scientific community.

## 1.7. BENEFITS

Implementation of the above recommendations would provide the U.S. scientific research community with a significant competitive advantage. Modernization of the nation's wide area networks by increasing speed, functionality, and size increases opportunities for research advances significantly.<sup>5, 6</sup> Greater network speed can reduce the time required to perform a given experiment and increase both the volume of data and the amount of detail that can be seen. Scientists accessing supercomputers would benefit particularly, because access speed is often critical in this work. Improved functionality frees scientists to concentrate directly on their experimental results rather than on operational details of the network. Increased network size extends these opportunities to tens of thousands of individuals located at smaller academic institutions throughout the nation. These modernization measures would significantly enhance the nation's competitive edge in scientific research.

The components of a shared network infrastructure would obviously benefit from global management, and the positive effects of such an approach would be widespread. Centralized administration of research in wide area networks would minimize duplication of effort and provide rapid resolution of identified high-priority problems. A global management structure would also allow a matrix approach to this distributed network expertise.

The U.S. communications industries would also gain a significant competitive advantage. Development of modern, low-cost distributed computing facilities for wide area networks would help maintain the United States position of world leadership in networking technology. Use of these products in support of science would accelerate the development of newer products by U.S. industry to meet challenges from both Europe and Japan. The United States would thus gain a position of world leadership in utilization of wide area, high bandwidth networks. This would increase the nation's competitive edge in communications technology as well as scientific research. As a spinoff, it would help maintain the U.S. leadership position in computer architectures, microprocessors, data management, software engineering, and innovative networking facilities.



## 1.8. WORKING GROUP PARTICIPANTS

Alison Brown  
Cornell University

David Stevens  
Lawrence Berkeley Laboratory

A. Fredrick Fath  
Boeing Computer Services

Bob Wilhelmson  
National Center for Supercomputing Applications

John Fitzgerald  
Lawrence Livermore National Laboratory

Philip Gross  
Mitre Corporation

Dennis Hall  
Lawrence Berkeley Laboratory

Jack Haverty  
BBN Communications Corp.

Charles M. Kennedy  
U.S. Army Ballistic Research Laboratory

Thomas Lasinski  
NASA Ames Research Center

Fred McClain  
San Diego Supercomputer Center

Patrick McGregor  
Contel Business Networks

Sandy Merola  
Lawrence Berkeley Laboratory

Hugh Montgomery  
Fermi National Accelerator Laboratory

Sushil G. Munshi  
United Telecom

Glenn Ricart  
University of Maryland

Richard T. Roca  
AT&T Bell Laboratory

Stan Ruttenberg  
University Corporation for Atmospheric Research

## 1.9. PRINCIPAL NETWORKING SITES

This list was compiled by the Lawrence Berkeley Laboratory, Informations and Computing Science Division.

Alabama	Anniston Army Depot Army Aeromedical Research Laboratory, Fort Rucker Army Missile Command Army Safety Center George C. Marshall Space Flight Center Gunter Air Force Station Maxwell Air Force Base University of Alabama, Birmingham University of Alabama, Huntsville	Lawrence Berkeley Laboratory (LBL) Lawrence Livermore National Laboratory (LLNL) Letterman Army Institute of Research Lockheed Palo Alto Research Laboratories Logicon, Inc. Los Angeles Air Force Station March Air Force Base Mare Island Mather Air Force Base McClellan Air Force Base McDonnell Douglas Computer Systems Company NAS North Island NASA Resident Office National Aeronautics and Space Administration Naval Air Station, Alameda Naval Ocean Systems Center Naval Personnel Research and Development Center Naval Post Graduate School Naval Sea Systems Naval Technical Training Center Naval Weapons Center Navy Elex Systems Engineering Center Navy Regional Automated Services Center Navy Supply Center Norton Air Force Base Presidio of San Francisco Rand Corporation Salk Institute San Diego State University San Diego Supercomputer Center Sandia National Laboratory - Livermore Schlumberger Caslab Science Applications Inc. - La Jolla Science Applications Inc. - Pleasanton Scripps Clinic and Research Foundation Scripps Institute of Oceanography Southwest Fisheries Center Stanford Research Institute International (SRI) Stanford Linear Accelerator Center Stanford University SUN Microsystems System Development Corporation Teknowledge, Incorporated Travis Air Force Base TRW Inc., Los Angeles University of California, Berkeley University of California, Davis University of California, Irvine University of California, San Diego University of California, San Francisco University of California, Santa Barbara University of California, Santa Cruz University of Southern California, Los Angeles University of Southern California, Marina Del Rey Vandenberg AFB Xerox Corporation
Alaska	University of Alaska, Anchorage University of Alaska, Fairbanks	
Arizona	Arizona State University, Tempe Davis Monthan Air Force Base Army Information Systems Command, Fort Huachuca Army Small Computer Engineering Center, Sierra Vista Army Yuma Proving Ground Kitt Peak Observatory Luke Air Force Base United States Geological Survey Astrogeology University of Arizona, Tucson	
Arkansas	Blytheville Air Force Base, Blytheville University of Arkansas, Fayetteville University of Arkansas, Little Rock University of Arkansas, Monticello University of Arkansas, Pine Bluff	
California	AMPC Advanced Computer Communications Advanced Decision Systems Aerospace Corporation Air Force Systems Command AMES Research Center Army DARCOM Logistic Control Activity Beale Air Force Base California Institute of Geology and Planetary Science California Institute of Technology Digital Equipment Corporation Dryden Flight Research Facility Eaton Corporation Edwards Air Force Base Electronic Data Systems Energy Applications & Systems Fleet Analysis Center FMC Corporation GA Technologies, Inc. George Air Force Base GTE TELENET Communication Corp. Headquarters, 6th Army, Presidio of San Francisco Institute for Advanced Studies IntelliCorp ITT/Federal Electric Corporation Jaycor Jet Propulsion Laboratory Kestrel Institute La Jolla Institute	
		Colorado JILA Lowry Air Force Base

	National Bureau of Standards	NEEACTPAC, Pearl Harbor
	National Center for Atmospheric Research	OINC NUWES, Hawaii Detachment
	Colorado State University	University of Hawaii
	Ford Aerospace and Communications	Wheeler Air Force Base
	Martin-Marietta Denver Aerospace	Idaho
	National Oceanic and Atmospheric Administration	Idaho National Engineering Laboratory
	Peterson Air Force Base	University of Idaho
	Science Applications Inc.	Illinois
	Solar Energy Research Institute (SERI)	Argonne National Laboratory
	University of Colorado	Astronomy & Astrophysics Center
	U.S. Air Force Academy	Chanute Air Force Base
	U.S. Army, Fort Carson	Fermi National Accelerator Laboratory
Connecticut		Gould Software Division
	Naval submarine School	Headquarters for AMCCOM
	Naval Underwater systems Center	Illinois Institute of Technology
	Yale University, New Haven	McDonnell-Douglas
Delaware		Naval Hospital
	University of Delaware	Northwestern University, Chicago
District of Columbia		Rock Island Arsenal
	Andrews Air Force Base	Scott Air Force Base
	Bolling Air Force Base	University of Chicago
	Defense communications Agency	University of Illinois
	Defense Mapping Agency	U.S. Army, Fort Sheridan
	George Washington University	Indiana
	NASA Headquarters	Grissom Air Force Base
	National Bureau of Standards	Indiana University
	National Science Foundation	Naval Weapons Support Center
	Naval Electronics Systems Security Engineering Center	Purdue University
	Naval Research Laboratory	University of Indiana, Bloomington
	Navy Regional Data Automation Center	University of Notre Dame, South Bend
	USAEASA	U.S. Army, Fort Ben Harrison
	U.S. Air Force, Pentagon	Iowa
	U.S. Department of Energy	Ames Laboratory
	U.S. Headquarters for the Department of the Army	Iowa State University
	Walter Reed Army Institute of Research	University of Iowa
Florida		Kansas
	Eglin Air Force Base	Kansas State University, Manhattan
	Fleet Training Center	McConnell Air Force Base
	Florida State University, Tallahassee	University of Kansas, Lawrence
	Homestead Air Force Base	U.S. Army, Fort Leavenworth
	Internet Systems Corporation	Kentucky
	Interscience, Inc.	University of Kentucky
	John F. Kennedy Space Center	U.S. Army, Fort Campbell
	MacDill Air Force Base	U.S. Army, Fort Knox
	Martin Marietta Corporation	Louisiana
	Naval Air Station	Barksdale Air Force Base
	Naval Coastal systems Center	England Air Force Base
	Naval Training Systems Center	Louisiana State University, Baton Rouge
	Patrick Air Force Base	Michoud Assembly Facility
	Service School Command	Navy Regional Data Automation Center
	Tyndall Air Force Base	Slidell Computer Complex
	University of Florida, Gainesville	U.S. Army, Fort Polk
	University of Miami	Maine
Georgia		Loring Air Force Base
	Auburn University, Auburn	Maryland
	Georgia Institute of Technology	Aberdeen proving Ground
	Marine Corps Logistics Base	Andrews Air Force Base
	Robins Air Force Base	David Taylor Naval Ship
	Skidaway Institute of Oceanography	Federal Data Corporation
	University of Georgia, Athens	Goddard Space Flight Center
	U.S. Army, Fort Gillem	Johns Hopkins University
	U.S. Army, Fort Gordon	NSSDC
	U.S. Army, Fort McPherson	National Bureau of Standards - Gaithersburg
	U.S. Army, Fort Stewart	National Computer Security Center
Hawaii		National Institutes of Health
	Camp H. M. Smith	National Security Agency, Fort George G. Meade
	Hickam Air Force Base	Naval Air Logistical Center

Naval Electronics Systems Command  
 Naval Surface Weapons Center  
 Network Solutions  
 Sea Automated Data Systems  
 Space Telescope Institute  
 University of Maryland  
 U.S. Army, Fort Detrick  
 U.S. Naval Academy

**Massachusetts**  
 Air Force Geophysics Laboratory  
 Atmospheric & Environmental Research, Inc.  
 Bolt Beranek and Newman Communications Corporation, Boston  
 Boston University  
 Computer Corporation of America  
 Digital Equipment Corporation  
 Dynamics Research Corporation  
 GTE Government Systems  
 Hanscom Air Force Base  
 Harvard College Observatory  
 Harvard University  
 Honeywell Information Systems  
 Intermetrics  
 Massachusetts Institute of Technology  
 MITRE Corporation  
 Northeastern University  
 Palladian Software, Inc.  
 Research Institute of Environmental Medicine  
 Thinking Machines Corporation  
 U.S. Army, Fort Devens  
 U.S. Army, Watertown  
 Woods Hole Oceanographic Institute

**Michigan**  
 Central Michigan University, Mount Pleasant  
 K. I. Sawyer Air Force Base  
 Marine Corps Air Station  
 Michigan State University, Lansing  
 Oakland  
 University of Michigan, Ann Arbor  
 Wayne State, Detroit  
 West Michigan University  
 Wurtsmith Air Force Base

**Minnesota**  
 Comten-NCR  
 Honeywell, Inc.  
 Sperry Corporation  
 University of Minnesota

**Mississippi**  
 National Space Technology Laboratories  
 Navy Norda

**Missouri**  
 Army Aviation Systems Command  
 Automated Logistics Management Systems Activity  
 Defense Mapping Agency  
 Marine Corps Central Data Processing Activity  
 McDonnell-Douglas Astronautics Corp.  
 University of Missouri, Columbia  
 University of Missouri, Kansas City  
 University of Missouri, Rolla  
 U.S. Army, Fort Leonard Wood  
 Washington University, St. Louis

**Montana**  
 Malstrom Air Force Base  
 Montana State University

**Nebraska**  
 Army Engineer Division  
 Offutt Air Force Base  
 Stanford Research Institute International (SRI)

University of Nebraska, Lincoln

**Nevada**  
 Nellis Air Force Base

**New Hampshire**  
 Dartmouth College  
 Frey Federal Systems  
 Pease Air Force Base  
 University of New Hampshire

**New Jersey**  
 Army Armament Research Development and Engineering  
 Army Armament Research  
 Army Communications  
 Army Information Systems  
 AT&T Bell Laboratories  
 Institute for Advanced Studies at Princeton University  
 John von Neuman Center  
 McGuire Air Force Base  
 Military Ocean Terminal  
 Princeton Plasma Physics Laboratory  
 Princeton University  
 Rutgers University, New Brunswick

**New Mexico**  
 Holloman Air Force Base  
 JSC White Sands Test Facility  
 Kirtland Air Force Base  
 Los Alamos National Laboratory  
 Sandia National Laboratories - Albuquerque  
 University of New Mexico  
 White Sands Missile Range

**New York**  
 Army Information Systems  
 Brookhaven National Laboratory  
 City University of New York  
 Clarkson College of Technology  
 Columbia University  
 Cornell University  
 European Command  
 General Electric Corp.  
 Goddard Institute for Space Studies  
 Griffiss Air Force Base  
 Grumman Aerospace  
 Hazeltine Corporation  
 Lamont-Doherty Oceanography  
 New York University  
 Plattsburg Air Force Base  
 Polytechnic Institute of New York  
 Rensselaer Polytechnic Institute  
 Rockefeller University  
 Rome Air Force Base  
 Sembach GE  
 State University of New York, Albany  
 State University of New York, Binghamton  
 State University of New York, Buffalo  
 State University of New York, Stony Brook  
 Syracuse University  
 University of Rochester  
 U.S. Military Academy

**North Carolina**  
 Duke University  
 Navy Regional Automated Services Center  
 North Carolina State University  
 Pope Air Force Base  
 Seymour Johnson Air Force Base  
 SRI Field Office  
 Triangle Universities Computation  
 University of North Carolina  
 U.S. Army, Fort Bragg

North Dakota	University of Tennessee Vanderbilt University
Minot Air Force Base, Minot	
North Dakota State University	Texas
University of North Dakota	Baylor College of Medicine
U.S. Air Force, Grand Forks	Baylor University
	Bergstrom Air Force Base
Ohio	Brooks Air Force Base
Case Western University	Carswell Air Force Base
Consolidated Data Center	Dyess Air Force Base
Defense Construction Supply Center	Geo-Chem Research Associates, Inc.
Defense Electronics Supply Center	Institute for Fusion Studies
Defense Systems Automation	Johnson Space Center
Lewis Research Center	Kelly Air Force Base
Newark Air Force Station	Lackland Air Force Base
Ohio State University	Laughlin Air Force Base
University of Cincinnati	Lunar and Planetary Institute
University of Toledo	Lyndon B. Johnson Space Center
Wright-Patterson Air Force Base	Microelectronics and Computer Technology Corporation
Oklahoma	Randolph Air Force Base
Headquarters, Department of Army, Fort Sill	Rice University
Oklahoma State University	Rockwell International
Tinker Air Force Base	Sheppard Air Force Base
University of Oklahoma	Southwest Research Institute
University of Tulsa	Texas A & M University
Oregon	Texas Accelerator Center
Intel Corporation	Texas Southern University
Oregon Graduate Center	University of Houston
Oregon State University	University of Texas
Portland State University	U.S. Army, Fort Bliss
University of Oregon	U.S. Army, Fort Sam Houston
Pennsylvania	Utah
Carnegie-Mellon University	Clearfield Federal Depot
Defense Mapping Agency	Dugway Proving Ground
Electronic Data Systems Network Service Center	Hill Air Force Base
Lehigh University	Tooele Army Depot
Leterkenny Army Depot	Utah State University
Millersville University	University of Utah
Naval Air Development Center	
Naval Aviation Supply Office	Vermont
Navy Ships Parts Control Center	University of Vermont
New Cumberland Army Depot	
Pennsylvania State University	Virginia
Philadelphia Naval Ship Yard	Army Materiel Command
Pittsburgh Supercomputer Center	BBN Communications Corporation
Systems Development Corp.	Boeing Computer Services
Temple University	CEBAF
Unisys Corporation	Center for Seismic Studies
University of Pennsylvania	College of William and Mary
University of Pittsburgh	Computer Sciences Corporation
Westinghouse Electric Corporation	Criminal Investigation Command
Rhode Island	Defense Advanced Research
Brown University	Defense Communications
Naval Data Automation Command	Defense Nuclear Agency
University of Rhode Island	Electronic Data Systems
South Carolina	Goddard Space Flight Center
Charleston Air Force Base	Honeywell Corporation
Clemson University	Langley Air Force Base
Myrtle Beach Air Force Base	Langley Research Center
NCR Corporation	Linkabit Corporation
Shaw Air Force Base	M/A-COM Government Systems
U.S. Army, Fort Jackson	Marine Corp Design Center
South Dakota	Naval Weapons Center
Ellsworth Air Force Base	Norfolk Naval Air Station
Tennessee	Science Applications Inc. - McLean
Maxima Corporation	Tamdem Computers, Inc.
Naval Air Technical Training Center	Teledyne Geotech Center for Seismic Studies
Oak Ridge National Laboratory	The MITRE Corporation
	U.S. Air Force, Pentagon
	U.S. Army, Fort Belvoir

U.S. Army, Fort Lewis  
Virginia Polytechnic Institute  
Wallops Flight Facility

Washington  
Battelle Northwest  
Flow Research Company  
Hanford Engineering Development Laboratory  
Stanford Research Institute International (SRI)  
Trident Training Facility  
University of Washington  
Washington State University

West Virginia  
University of West Virginia

Wisconsin  
University of Wisconsin, Madison  
University of Wisconsin, Milwaukee

Wyoming  
F.E. Warren Air Force Base

## 1.10. REFERENCES

1. S. Ruttenberg, "Summary of Agency Network Plans," *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States (June 1987)*.
2. D. Stevens, "DOE Networking Requirements," *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States (June 1987)*.
3. A.F. Fath, "Networking Trends in Data Communications," *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States (June 1987)*.
4. G. Ricart, "Research University Bandwidths," *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States (June 1987)*.
5. D. Hall, W. Johnston, Hutchinson M., et al, "The Effect of Distributed Computing Technology on Wide Area Network Capacity Requirements," *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States (June 1987)*.
6. T.A. Lasinski, "Role of Graphics Workstations in Supercomputing," *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States (June 1987)*.
7. J. Haverty, "Networking: Some Observations on the Future," *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States (June 1987)*.
8. S.G. Munshi, "Industry and Technology Trends," *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States (June 1987)*.
9. R.T. Roca, "AT&T Presentation Graphics," *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States (June 1987)*.
10. P. McGregor, "Contel Presentation Graphics," *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States (June 1987)*.
11. "The Future of Fiber", *Telephone Engineering and Management*, p.52 (February 1, 1987).

## 2. SPECIAL REQUIREMENTS FOR SUPERCOMPUTER NETWORKS

---

Dr. Robert Borchers  
Lawrence Livermore National Laboratory

### Abstract

*The goal of our society's investment in supercomputers is to retain and enhance U.S. economic and scientific competitiveness and national security through enhanced productivity of our scientific researchers. To assure continued attainment of these goals, it is necessary to provide state-of-the-art tools to the U.S. scientific research community. The size of the supercomputer investment, including only the sites represented by participants in this working group, exceeds \$0.5B capital investment and several \$100M/year annual operating expenses spread across the Department of Defense, Department of Energy, Department of Health and Human Services, National Aeronautics and Space Administration, and the National Science Foundation. Continued provision of high-quality, remote access to these facilities is necessary for the continued productive use of this investment. Continued remote access to these facilities is required to assure continued U.S. leadership in science.*

*The most successful remote access to these facilities has been achieved where the establishment of the required remote access mechanisms is left to those with the responsibility to provide the supercomputer resources. This has led to the establishment of networks (MFENET, NASnet, SDSCnet, etc.), with particular topology, connectivity, and functionality optimized to serving the intended community of researchers. Other network functions, such as electronic mail, have come along as ancillary benefits. Now we are challenged to supply supercomputer access to a much larger community of researchers dispersed by geography and discipline. To achieve this goal, we must not discard the techniques and priorities that have led to our successes to date. We must build on and expand the existing environment.*

*In the supercomputer environment, performance (involving number of hops, priority level of message service, priority routing, and reliability, as well as bandwidth) is paramount. Connectivity and other considerations are secondary. This focus must be retained to assure continued success in providing supercomputer access to the scientific community.*

### 2.1. REQUIREMENTS

Supercomputers are, by definition, the most powerful computers available at any given time. Today's supercomputers are capable of performing billions of floating point operations per second and generating billions of bytes of data in memory. In order to exploit the power available from supercomputers, networks have been created to provide remote access and file transfer capabilities to users. An important secondary benefit of these networks has been to enhance user cooperation and collaboration through electronic mail and other network facilities.

Future supercomputer networks must continue to be driven by these same requirements--remote access and file transfer. Of course, the exponential growth in data production rates and in the ability of users to productively digest these data places enormous demands on future networks. The failure to provide networks that adequately match these data production and consumption rates will seriously inhibit scientific productivity and the entire investment strategy in supercomputers.

Generic requirements for an adequate network to service supercomputer users include negligible delay times, high peak bandwidths available on demand and adequate continuous bandwidth, and transparent operation to the user.

The large volumes of data produced by supercomputers routinely require graphical presentation to the user. In order to obtain the maximum scientific benefit, the graphics presentations should be of high resolution, interactive with the user, and available in adequate time sequences. It is imperative that the network serving the supercomputer be able to deliver these services.

The supercomputer network must be able to provide service to geographically concentrated groups of users that require a relatively continuous and large amount of service, as well as smaller, more distributed groups that may appear or disappear over shorter periods of time.

To successfully build a national supercomputer network to support scientific research, the following issues, at a minimum, require consideration. Of highest priority is performance. Whatever configuration is proposed must be first accessible, continuously available, serviceable, and reliable. Its bandwidth must accommodate the varying requirements of the users and not diminish effectiveness as the user base grows. The system must sustain through time a minimum service level commitment to all users. This commitment is sustained through an authority/responsibility infrastructure that considers effective and efficient management control and funding mechanisms. And, finally, the system must dictate minimum standards and security requirements.

1. **Control:** Managers of supercomputer centers require control of the access mechanisms to ensure quality service to their users.
2. **Accessibility/Connectivity:**  
User to supercomputer--highest priority  
Supercomputer to supercomputer--second highest priority  
User to user--lower priority  
Currently available access methods take considerable time to establish.
3. **Bandwidth:** There is a spectrum of services currently required. On one hand, there is the need for very high (100-Mbyte/s) bandwidth for a small number of users with highly interactive graphic requirements. On the other hand, there is a need for a large number of users to have relatively low (9600-baud) bandwidth access. During the next 15 years, both of these bandwidth requirements will increase substantially. Currently, the high bandwidth requirement is not commercially available over large distances.
4. **Reliability and Recovery:** Access to supercomputers must achieve 99.9% availability through redundancy or automatic recovery. Mean time to repair should not exceed a few minutes. Long interruptions in access to supercomputers are likely to idle a very scarce and valuable resource.
5. **Standards:** The purpose of standards is to allow one to purchase off-the-shelf equipment that can be plugged together.
6. **Security/Privacy:** Protocols and hardware interfaces must take into account the need for security/privacy.



7. **Topology:** Performance areas such as bandwidth, connectivity, and reliability must be taken into consideration.
8. **Responsiveness:** Interactive use of supercomputers (e.g., symbolic/dynamic debugging) requires that the network be responsive to interactive requests.
9. **Management:** A 24-hour operational environment requires network monitoring that responds to service interruptions in a timely manner.
10. **Cost:** A network should allow management to determine where resources are being expended and to determine value rendered.

## 2.2. STATUS

We will break existing networks roughly into three categories and examine how each has dealt with the three issues (performance, access, and control) and to what results. Finally, we make some observations about existing centers and which of these categories they use. The categorization of networks is done along the dimension of

- high performance, smaller user community, and targeted types of service, and
- low performance, store and forward, very wide spread, and mail service.

### 2.2.1. Category I: Mission-Oriented Networks

These types of networks include the National Aeronautics and Space Administration (NASA) Numerical Aerodynamics Simulation (NAS) program, the National Magnetic Fusion Energy Computing Center (NMFEC), and the San Diego Supercomputer Center (SDSC).

- These centers see themselves as networks providing client-center service for supercomputer applications. This allows design choices that markedly enhance performance.
- These centers have well-defined clientele and are "comparatively closed."
- Being mission and center oriented, these centers retain control that allows design choices to be made and facilities management, fault isolation, user services, and repairs to be enhanced.

We note that NAS has conducted their network within an open architecture.

### 2.2.2. Category II: Confused Orientation

This category serves all research institutions, although the original concept was similar to Category I. The NSF/ARPA internets are examples of Category II orientation.

- Performance has *not* been aided by focusing on the type of service to be provided or by focusing on who is being served. The Defense Communications Agency (DCA) is, however, attempting to deal with general load problems. Increased usage is rendering this internet "barely functioning" from time to time. It is not clear whether current attempts to overcome the diffuse goals and control problems will succeed.
- These nets have a mandate to cover an indefinite number of institutions. Connection is *not* being limited to small numbers. Although the backbone started out as "center oriented" and although many describe its goal as "client to center," it is following a "client-to-client" style unlikely to be able to address supercomputer access needs.
- Control is diffused between DCA/NNSC/regional/campus/backbone sites. It remains to be seen whether the current interim management structure can work in a coherent fashion and also serve the needs of the supercomputer centers.

### 2.2.3. Category III: Serving all Comers

This category is very wide spread and contains minimum functionality (e.g., BITNET, CSNET).

- "Performance" as such is *not* the issue other than that "the mail will usually get through ..."  
When used for more demanding applications (e.g., remote job entry), results are mixed.
- Access is very wide and no real attempt is made to limit access, especially in pursuit of performance!
- Orientation is "flat" and not "towards a central facility." Therefore, control is vested in a separate organization (and on campuses may be very diffuse to nonexistent), and is not coordinated with any activity outside of the net itself.

### 2.2.4. Category IV: Ad Hoc Solutions

This category includes the use of dialup facilities at 9600, or 2400, or 1200 baud, or via X.25. This is not really a network but rather a way to spread access into nets and into centers.

### 2.2.5. Summary of Category Use

- The supercomputer centers vary in their use across these categories.
- Mission-oriented centers use mostly Category I.
- NSF centers are currently using all categories with differing "centroids."
- SDSC is mostly Category I,
- Pittsburgh Supercomputer Center is spread across Categories II, III, and IV,
- Cornell is mostly Category IV and is trying to migrate towards Category II.

In summary, the ability to target and trade off and optimize both for a client-center orientation and for type of service to be provided seems to be *very* important for serving the needs of supercomputer center users. Hopefully, these advantages can be retained as the spread of the net increases. We also hope to see this done within an open architecture.

## 2.3. CONCLUSIONS

Supercomputer access can be provided remotely and successfully by networks. Performance, as viewed by the scientists using the supercomputers, is the key criterion for a successful network. Responsiveness, accessibility, and reliability are primary considerations. The preceding sections discussed the issues related to providing good performance and listed examples where networks are supporting large remote user communities.

On the other hand, networks designed primarily for "connectivity" have been less successful. This leads us to the conclusion that provision of supercomputing services places a more severe constraint on network design and implementation. The organization that has the responsibility for providing supercomputing resources must have the authority to see that the network is properly designed, implemented, and managed.

Currently, access to the new supercomputer centers is still very spotty for many users. In assessing an overall network strategy for supercomputing, these problems will take some time to resolve. Centers (many already in full operation) will need to take interim steps to provide remote user access.

## 2.4. WORKING GROUP PARTICIPANTS

F. Ron Bailey  
NASA Ames Research Center

Robert R. Borchers  
Lawrence Livermore National Laboratory

Charles Crum  
National Cancer Institute

Dennis Duke  
SCRI-Florida State University

Dieter Fuss  
Lawrence Livermore National Laboratory

Sid Karin  
San Diego Supercomputer Center

James Leighton  
LLNL/National Magnetic Fusion Energy Center

Michael Levine  
Pittsburgh Supercomputing Center

Norman Morse  
Los Alamos National Laboratory

Ari Ollikainen  
NASA Ames Research Center

Harry L. Reed, Jr.  
Ballistic Research Laboratory

Rudi F. Saenger  
Naval Research Laboratory



### 3. INTERNET CONCEPTS

---

**Dr. Lawrence H. Landweber**  
University of Wisconsin, Madison

#### **Abstract**

*The Working Group on Internet Concepts recommends that it should be a national priority to provide a ubiquitous computer communication capability to interlink the entire U.S. academic community by the year 1992. This should include a minimal basic service level for all participants as well as the possibility of a higher level of service when needed to support a particular application. This national communications infrastructure can be based on existing agency and academic networking activities. To be successful, significant new funding will be required.*

*After a period of steady evolution of networking technology, we are beginning to experience a discontinuity that should cause us to start thinking now about possibly changing the basis upon which research networks of the future will be built. This discontinuity is the unexpectedly rapid growth of terrestrial capacity due to the widespread installation of optical fibers. Already this technology is beginning to have visible consequences in the design of local area, metropolitan area, and wide area networks.*

*Research and planning should begin now so that research networks can be based on the emerging technologies described as these technologies become widely available.*

#### **3.1. SUMMARY OF CURRENT RESEARCH NETWORKING IN THE U.S.**

A number of research networking facilities are currently under development or in operation in the U.S. under the support of various government organizations, including the Departments of Defense and Energy (DoD and DOE), the National Science Foundation (NSF), and the National Aeronautics and Space Administration (NASA). These research network facilities are summarized below.

##### **3.1.1. DoD Networks**

The U.S. DoD, largely through the Defense Advanced Research Projects Agency (DARPA), has contributed to the development of U.S. computer communication technology for the last two decades. The ARPANET, the world's first national packet switching network, was developed in the late 1960s to support resource sharing among some 30-50 DARPA-sponsored research groups. This networking technology has continued to evolve and expand, and is now used to support other government-sponsored academic research activities. The packet switching technology has been applied to other communication media such as mobile radio (PACKET RADIO), broadcast satellite (SATNET, WBNET), and local area networks. Also, a set of standard communication protocols and gateways has been developed to permit the interconnection of various packet nets into a loosely coupled internet. DARPA and the Defense Communications Agency continue to operate this internet and its constituent subnetworks and gateways as a major element of the U.S. defense research infrastructure. Protocols developed under the DARPA program were adopted by the DoD as a whole, and are used widely in other research networks as described below.

### 3.1.2. NSF Networks

#### 3.1.2.1. CSNET Description

CSNET, the Computer Science Research Network, was funded by NSF in 1981. CSNET's goal was to develop a network that would serve all U.S. computer researchers in universities, industry, and government. CSNET was initially funded by NSF for 5 years. It is now fully supported by its approximately 190-member organizations and no longer receives a subsidy from NSF.

CSNET uses a tiered approach, with different service levels available to groups depending on their needs and ability to pay. The Phonetnet component involves telephone-based relaying of electronic mail to and from a central mail server, which is also connected to the ARPANET. The second component, X.25NET, involves use of X.25 public data networks. The Transmission Control Protocol/Internet Protocol (TCP/IP) is run over X.25 and an IP gateway is provided to enable interoperability with other TCP/IP academic/research networks. The third CSNET component uses ARPANET. Most university computer science departments with ARPANET connections are members of CSNET.

CSNET is managed by the University Corporation for Atmospheric Research (UCAR). Technical support and user services are provided by a Coordination and Information Center at Bolt, Beranek and Newman under contract to UCAR.

#### 3.1.2.2. NSFnet Description

NSFnet is intended to support resource sharing and research collaboration within the national scientific research community. This sharing requires a high bandwidth, reliable communication system that spans the U.S. geographically, interconnecting university, federal, and private research laboratories and the resource centers. Among the major resources to be shared through NSFnet are the national supercomputer centers; other examples include major library facilities and large-scale experimental facilities.

NSFnet has been structured as a three-level hierarchy: local area networks on the campus and laboratory sites; regional, state, consortium, and discipline-oriented networks tying the campuses together; and a "backbone" network spanning the U.S. to provide national connectivity.

Construction of NSFnet is progressing simultaneously in all three levels. As local area networks are proliferating rapidly on university campuses, they are being organized into local campus networks. Over a dozen middle-level networks are in the process of being organized. Five of these are sufficiently advanced to be supporting user communication, and the others are making progress. Finally, an initial subset of the national backbone, with hubs at NSF-supported supercomputer sites, is fully operational.

To meet the NSFnet objectives, the backbone network will need to be upgraded substantially with more nodes, higher bandwidth (multiple DS1 by 1992), and comparatively faster packet switches. Considerable engineering and development on routing algorithms and protocols and on operating and monitoring capabilities will be needed for this to be successful.

The regional networks will need the same technical developments. In addition, they have a large task of providing user workshops and information services to the constituent campuses. One of the significant benefits of NSFnet's three-level hierarchy will be the widespread dissemination of computer

networking knowledge and understanding within the university community; a more centralized organization would not have this benefit. This large and important training task will fall largely on the regional networks.

### 3.1.3. DOE Networks

#### 3.1.3.1. MFENET/ESnet

The major role of MFENET is to provide nation-wide access for over 4,000 users of the National Magnetic Fusion Energy Computer Center (NMFEECC). The major supercomputers at the center include two CRAY-1s, a CRAY X-MP/22, and a CRAY-2. Approximately 60 other host computers are now connected to the MFENET. The MFENET is currently undergoing a major redesign to support several new requirements, including supporting much higher bandwidth in the future, providing a migration path for support of the International Standards Organization (ISO) standard protocols, and allowing participation in interagency internetworking efforts. The new design has also been chosen as the basis for an "internet" project within DOE, which has been designated as the Energy Sciences Network (ESnet). ESnet is intended to support the networking requirements of the programs supported by the Office of Energy Research (OER) within DOE. The goal is to enhance the networking capabilities of these programs and, at the same time, to reduce costs associated with redundant or incompatible data communications projects. Because of the international scope of several of the programs supported, network connections to several foreign sites are in progress. Sites of interest are in Switzerland, West Germany, and Japan. Additional interest exists for several other nations.

#### 3.1.3.2. HEPnet/LEP3NET

HEPnet and LEP3NET are networks that support the High Energy Physics (HEP) programs of the DOE and NSF. HEPnet is currently a large international network of Digital Equipment Corporation (DEC) computers, and is in the process of being restructured in the U.S. into an X.25-based multivendor network. LEP3NET is an X.25-based leased line network that came online in 1985 to serve the programmatic needs of L3, the largest HEP experiment. It also serves other sectors of HEP to a limited extent. Its X.25 architecture, which was mandated by the international character of the network, has become the near-future model for HEPnet. HEP is a strongly collaborative program that involves large experimental collaborations as well as some theoretical projects. Sharing of databases, cooperative development of computer software bases for data analysis, remote monitoring of experiments, and the distribution and sharing of experimental data between laboratory sites and home institutions of the physicists has made computer networking a necessity for the functioning of the field.

The increasing importance of networks, and the need to reorganize HEPnet into a more structured, managed, and multivendor entity, has been driven by the nature of the next generation of HEP experiments. The leading HEP collaborators, who are building experiments at CERN (Geneva), Fermilab (Chicago), SLAC (Palo Alto), and other HEP laboratories, now include 100 - 400 physicists, from 10 - 40 institutions, in as many as a dozen countries each. Over the course of one experiment (over a period of several years),  $10^{13}$  bytes (100 Tbits) of data may be generated and stored for analysis and reanalysis. These data are already highly compressed, formatted, and partially analyzed at the time of recording.

The bandwidth needed today per HEP group (6 - 10 physicists) for full sharing of programs and databases is in the 56-kbit/s range, but is limited by economics to the 9.6- to 19.2-kbit/s range. Sharing and distribution of data files between the laboratory sites and home institutions, which would be of enormous benefit to U.S. physicists even today, will only be possible when bandwidths of 1.5 Mbit/s are affordable (or are provided by national networks) to individual physics groups.

In the recent past, HEPnet arose as a "grass roots" effort, driven by the needs of HEP experimenters. The predominance of DEC computers at experimenters' home sites led to the formation of a homogeneous DECnet, which was soon extended nationwide. The network then came in contact with NASA's SPAN DECnet at the California Institute of Technology and with the European HEP DECnet through the LEP3NET transatlantic link. This brought it to its current size of more than 1000 nodes.

HEPnet is now being rebuilt in the U.S. around a 56-kbit/s backbone that will connect the Lawrence Berkeley Laboratory (LBL), Fermilab, Brookhaven, and the Massachusetts Institute of Technology (MIT). International links include LBL-Japan and Fermilab-CERN (a 64-kbit/s satellite link). The network will be X.25 based, with X.25 switches at the primary node sites programmed and managed by the HEP laboratory staff. The number of feeder circuits to the backbone in the U.S. is expected to grow to approximately 60, at a typical speed of 9.6 kbit/s.

HEPnet will run both DECnet and the Coloured Books protocols over X.25. The Coloured Books (used extensively in the U.K. academic community over the last several years) will allow file transfer between DEC and IBM computers. This method of operation has been used effectively on the international LEP3NET. It has been chosen because of the X.25 requirement for European connections and the fact that a network of this type may be implemented inexpensively using standard commercial products.

HEPnet will also provide access to computers and to Data PBX switches at the HEP laboratories. This is the means of terminal access to the HEP laboratory's central computing facilities.

In the longer term, HEPnet will migrate towards use of the ISO protocols as an integral part of the ESnet. Once the ESnet gateways are ready, the HEPnet X.25 traffic may be switched over ESnet lines. In preparation for the ISO transition, HEPnet will also begin use of IP-based protocols, in parallel with its existing services, in the near future.

It should be noted that HEP networking is, in contrast with much of the research carried out over some other networks, a group- and mission-oriented, full-time activity. Currently a strong need exists for full-time access to remote facilities, which would expand in step with computer communications and HEP detector technology, if there were not strict funding limitations.

By 1992, given the foreseeable needs of the large experiments, each physics group could fully and effectively utilize at least 1.5 Mbit/s for its research needs. This does not include any transmission of non-numerical "low information density" data such as screens of pixels.

By the year 2000, full participation of each physics group in the activities of its remote experiment, at the Superconducting Supercollider for example, could entail full-time use of a link in the range of 10 - 100 Mbit/s.



### 3.1.4. NASA Networks

NASA operates three types of communication systems for its scientific programs: mission support communications (MSC) for the command, control, data collection, data distribution, and data analysis in real time; program support communications (PSC) for administrative uses including the computer-to-computer and terminal-to-computer scientific activity not in real time; and center or facility communications that link facilities and individual investigations to either or both MSC and PSC systems. Neither the MSC nor the PSC systems provide computer networks with interoperability or full services (remote login, electronic mail, file transfer, routing) except in very limited, specialized projects or programs as addressed below. Circuit networks are actually collections of point-to-point links, custom interfaces, and heterogeneous equipment that are centrally managed and maintained by Goddard Space Flight Center (GSFC) for MSC or Marshall Space Flight Center (MSFC) for PSC.

At the centers and facilities, serious systematic efforts have been undertaken to share computing resources using computer networks and to extend these computer networks beyond the center or facility geographic limits to support missions or programs that involve other centers, facilities, universities, federal laboratories, and industrial sites.

Development of computer networks, as compared with circuit networks that provide interoperability between computing systems, has been a strong, rapidly growing "grass roots" effort by administrators, managers, and scientists alike. NASA joined the first computer network, the ARPANET, at Ames Research Center in 1970 as part of a joint agreement between NASA and DARPA to share a jointly developed supercomputer facility, the ILLIAC IV, with an extended group of computational physicists distributed geographically. Subsequently, GSFC and the Jet Propulsion Laboratory (JPL) have linked to the Defense Data Network (DDN, ARPANET, MILNET) to meet some programmatic requirements.

The Space Physics Analysis Network (SPAN) is an important "grass roots" effort to build a computer network using nearly homogeneous computer systems and vendor-specific communications protocols overlaid on circuits provided by the PSC system. SPAN began in 1980 from recommendations of a NASA/university/industry space science user committee (called the Data Systems Users Working Group, DSUWG) and became operational in 1981. SPAN uses the DECnet protocols and currently (as of February 1987) contains approximately 650 space and earth science research computers and is internetworked transparently to nearly 1000 others (from cooperating networks such as the HEPnet and the Texas Academic Network, TEXnet). SPAN uses 56-kbit/s links between four routing centers (GSFC, MSFC, Johnson Space Flight Center, and JPL). All other institutions are connected to the nearest SPAN routing center at 9.6 kbit/s. In addition to the internetwork of DECnet networks, SPAN has several gateways onto ARPANET and BITNET. The DSUWG remains today the SPAN advisory group and recommends that the DECnet protocol be replaced with the ISO/OSI standard as it becomes available.

The advocacy and startup of the Numerical Aerodynamic Simulation (NAS) program, managed at ARC, requires nationwide access to this national facility at data rates in excess of anything available through the DDN or public packet networks (i.e., 1540 kbit/s and up). To support the unique design of the NAS facility, a star network has been implemented utilizing the new technology of Ethernet long-haul bridges (Vitalink boxes) overlaid on the agency-wide PSC backbone (multiple T1 trunks) and tail circuits from non-NASA nodes to the remote NAS user. DoD standard protocols (TCP/IP and application suite) are employed. The circuits are provided by the PSC system. The computer network, NASnet, is implemented by ARC specifically to meet the high demand for response time critical communications of a high-performance remote work station linked to an ensemble of supercomputers.

The growth of costs resulting from the demands for more and more bandwidth and for more and more connectivity to science practitioners, facilities, and support activities by such networks as SPAN and NASnet; the diffusion of multiple vendor computing products throughout the scientific infrastructure; the profusion of dissimilar computer networks at various NASA centers; and the growing need for interconnecting NASA networks with other networks (ARPANET, MILNET, NSFnet, BITNET, SPAN, etc.) have focused the need for centralized coordination of data communications for computer networks for science. The NASA Science internet (NSI) is a full-service, packet switched network that will provide full services utilizing the standard DoD standard protocols, Internet Protocol (IP) and Transmission Control Protocol (TCP). NSI will link the science communities together electronically, including the disciplines of astrophysics, astronomy, and land, ocean, climate, and planetary sciences, as well as tie these communities to other interoperable networks overlaid on PSC system-provided circuits. By aggregating science data communication requirements and supporting them with a computer network, moderation of cost growth in circuits will be achieved. Further, the connectivity of scientist to computing resources is increased. The NSI will make the transition from the DoD protocols to the ISO/OSI protocols in a phased, controlled manner that meets the science networks requirements and accommodates other protocols such as DECnet where technically feasible over the same physical circuits.

## 3.2. THE ACADEMIC INTERNET OF 1992

### 3.2.1. Major Conclusions

The 1992 aspirations in this report are based largely on existing technology or evolutionary changes from current technology. There are already known or foreseeable limitations to currently available networking technology that will have to be addressed in the short and long term.

The major conclusion for 1992 is as follows:

It should be a national priority to provide a ubiquitous computer communication capability to interlink the entire U.S. academic community. This should include a minimal basic service level for all as well as the possibility of a higher level of service when needed to support a particular application. This national communications infrastructure can be based on existing agency and academic networking activities. To be successful, significant new funding will be required.

This goal is motivated by the growth of special national research resources such as the supercomputer centers; computer-based libraries and database resources (such as the National Library of Medicine, Genetics Bank, and other biomedical information resources); and major experimental facilities such as the national physical laboratories and existing and planned space science facilities (including the space station).

A computer networking structure for the entire academic research community would vastly enhance access to research results and stimulate the exchange of information through the use of electronic messaging, online electronic publishing and archiving, and remote access to computing and database resources. Such facilities would make possible the coordination and conduct of experiments and research requiring special national resources otherwise impossible today. The U.S. academic system is one of its most important national resources and is the primary source of new technology. Facilitating access to and the exchange of knowledge are essential elements of maintaining the international competitive advantage for the U.S.

### 3.2.2. Technology Issues for the 1992 Research Internet

In the 1992 time frame there will be roughly 1000 campuses supporting 1,000,000 researchers. A large number of these campuses will have local area networks running at 10 Mbit/s, although some larger campuses will have 100-Mbit/s local area network backbones. These campuses will be interconnected by long-haul networks organized in a variety of ways. Agency and discipline networks will serve the needs of user communities with a common orientation or mission. Consortia networks will provide for efficient access to shared resources, such as supercomputers, by their biggest users. Geographic networks will provide for the general interconnection of campuses or users without special needs or affiliation. One or more backbone networks will carry transit traffic between geographic networks.

The applications for this network will require a minimum of 64 kbit/s to user workstations. Some applications, such as workstations interacting with supercomputers, may require high speeds such as 1.5 Mbit/s or greater for one user. In addition, groups of experimental researchers will, before 1992, require upwards of 1.5 Mbit/s for data file transfer and for sharing of programs and databases between the computing facilities at their home institutions. One implication of these data rates is that each campus will need at least 1.5-Mbit/s capacity into the internet, since this bandwidth must be shared between campus users.

Backbone circuits, or other high usage circuits in the internet will require speeds in excess of 1.5 Mbit/s. Facilities of 45 Mbit/s may be required in some instances. Available Time Division Multiplexing (TDM) technology will make it practical for 45-Mbit/s service to be shared between agencies and between voice and data.

We foresee several technological problems in achieving this environment. These problems relate to its scale, scope, and complexity. This environment is more than one order of magnitude larger than existing computer networks or internetworks in terms of the number of users, data rates, etc.

The first problem is availability of high-speed packet switches. The majority of switches available today support data rates of 64 kbit/s, with a few high-performance switches supporting a small number of 1.5-Mbit/s circuits. The 1992 internet will need packet switches capable of operating at 1.5-45 Mbit/s. The hardware and software technology necessary to build these switches exists today, but product development lead times and uncertain markets for such switches may impede their availability.

The second problem is routing of packets in this internet. Existing routing algorithms for large networks are based on hierarchical topologies. It may not be possible to impose a hierarchical structure on the 1992 internet. In addition, constraints on routing must be imposed to reflect organizational realities. Mission-oriented networks need to continue to meet their mission-oriented needs and must be protected from excess transit traffic. Existing routing algorithm technology does not fully address these needs.

The third problem is one of robustness. Firewalls must be constructed to prevent hardware or software failures from propagating. Without these firewalls, each increase in scale will result in a decrease in internet availability. In addition to concern over hardware and software failures, some attention must be given to securing the network from malicious users.

The fourth problem is one of performance. As recent experience has shown, the performance of internetworks constructed out of overloaded networks and overloaded gateways can be poor. The scale of the 1992 internetwork will aggravate these problems. These problems could be alleviated by the

installation of backbones of sufficient capacity between strategic points of the internet, following network engineering studies to determine the optimum backbone locations. In addition, more technology is needed in controlling congestion and providing efficient and fair allocation of network bandwidth within desired policy guidelines.

### 3.2.3. Opportunities for Resource Sharing and Coordination

The various government networking activities touch a significant segment of the U.S. academic research community. The interlinking of some of these networks has already begun (e.g., NSFnet, the regional networks, the supercomputer networks, ARPANET, and other experimental defense nets). Moreover, most of these networks are adopting the TCP/IP protocol suite in the short term to take advantage of existing, widely tested, available implementations and experience already gained in operating the internet of networks using these protocols.

This commonality among the major government research nets creates an opportunity for the sharing of resources and coordination of development efforts required to accommodate the increasing scale of the internet and to coordinate planning for the new technologies needed to meet foreseeable research computing networking requirements of the 21st century.

The working group makes the following recommendations concerning research networking efforts at NASA, DoD, DOE, NSF, the National Institute of Health (NIH), and the National Bureau of Standards (NBS):

1. These organizations should be permitted and encouraged to seek ways of sharing resources, for example, through pooling of long-haul transmission requirements and multiplexed use of common 45-Mbit/s long-haul channels. Consideration should be given to the use of digital cross-connect technology to permit flexible allocation of capacity to meet various requirements.
2. Joint engineering efforts should be mounted to support the modeling, measurement, and management of the research internet. A capability for internet capacity and topology planning should be developed.
3. Joint R&D efforts should be initiated to address the problems of routing, congestion control, robustness, and performance in the internet, and efforts to couple the results into migration plans from the DoD to the ISO protocols should be initiated.
4. User services and information clearinghouse functions should be coordinated and, where feasible, combined to provide maximum benefit to the entire research community.
5. Development and operation of host/mail Name Servers and routers should be coordinated.
6. A pilot program to procure (develop?) 45-Mbit/s packet switches should be put in place in anticipation of post-1992 needs.

### 3.2.4. Summary

The working group strongly endorses the cooperative sharing of resources by the government organizations supporting research networks and urges a national initiative to support networking across the entire U.S. academic community.

### 3.3. ACADEMIC NETWORKING IN THE YEAR 2000

This section represents some ideas about how user requirements and the technology available to meet those requirements will evolve between now (1987) and 2000. It will be argued that, after a period of steady evolution of networking technology, we are beginning to experience a discontinuity that should cause us to start thinking now about possibly changing the basis upon which research networks of the future will be built. This discontinuity is the unexpectedly rapid growth of terrestrial capacity due to the widespread installation of optical fiber. Already this technology is beginning to have visible consequences in the design of local area, metropolitan area, and wide area networks.

#### 3.3.1. User Requirements

##### 3.3.1.1. Overall Number of Users

The community of users who will need to be networked together in 2000 was estimated to be about  $10^6$ , distributed over some  $10^3$  campuses, national laboratories, and other research facilities. At any given instant of time, it was estimated that an average of about 500 users per location would be active on the network in the sense that their traffic would actually be flowing across the system.

##### 3.3.1.2. Classes of Traffic

Six classes of traffic were defined, which we felt captured most of the sort of activity we could predict would take place. One class that was deliberately excluded is voice telephony, on the grounds that these resources will continue to be procured by other noneducation/research departments. This is because voice telephony is based on completely different bit rate, line quality, and other requirements, and is of such a volume as to completely distort the design point of any reasonable research computer network. If, between now and 2000 the telephone companies or others make great progress in designing and installing networks, all of whose nodes carry mixed voice/data traffic, this omission would have to be revisited, but we expect progress toward this objective to be slow. In the meantime, even though integrated **switching** may not be widely installed, integrated **transmission** will be the norm, as it is today (e.g., combining data and voice on local access lines to the central office).

- Class 1: Video and teleconferencing. Result: 50-Mbit/s peak bit rate, 1 Mbit/s average.
- Class 2: Closely coupled "supercomputers". Examples might be a Connection Machine and a Cray, each executing different steps of a joint program job. Result: 3-Gbit/s burst rate.
- Class 3: Workstation-to-host, closely coupled. The workstation user must interact with the supercomputer's execution stream in real time. To do this the user must see one  $10^6$  pel frame every 30th of a second. Result: 300-Mbit/s peak bit rate.
- Class 4: Workstation-to-host, loosely coupled. One  $10^6$  pel frame per second. Result: 10-Mbit/s peak bit rate.
- Class 5: Fast packet switching. Result: 1-Mbit/s peak bit rate.
- Class 6: Everything else. This especially includes packet switching as we now understand and use it. This also includes incidental voice for order-wire purposes, fax, modem traffic, and 64-kbit/s leased or circuit switched data. Result: 64-kbit/s peak bit rate.

### 3.3.1.3. Required Bit Rate

The following table gives our estimate of typical numbers of concurrent users who might be active within each traffic class, and the total aggregate bit rate generated by these users. This traffic will be widely dispersed geographically, but it is nevertheless instructive to get a network-wide estimate of total ongoing traffic requirements.

Class	No. Simultaneous	Bit Rate (Gbit/s)	Required Switchover Time	Required Frequency of Switch
1	100	5	Seconds	Minutes
2	100	300	Seconds	Hours
3	1000	300	Seconds	Hours
4	10,000	10	msecs	Seconds
5	10,000	10	Sub msec	msec
6	10,000,000	64	msecs	msec

Adding up the total required bit rate gives 779 Gbit/s, which we can round off to 1000 (1 Tbit/s). Again, this is the aggregated bit rate distributed geographically over the network, which is perhaps a highly artificial way of counting bandwidth but is at least a way of giving an approximate upper bound on demand.

### 3.3.1.4. Switching

The last two columns are an attempt to capture the requirement on the switching technology of the future. The next to last column gives estimates of the time the user requires his/her path to be switched from one direction to another, and the last column gives an estimate of how often it will typically be necessary to do so.

### 3.3.1.5. Inhomogeneity

One final requirement, the ability to accommodate inhomogeneity, was also discussed. It was our feeling that the research community has particularly strong needs for its members to use systems whose characteristics are not forced into the same mold just so they can become parts of a communicating community. To do so would constrain the very innovation that the community is attempting to achieve. The present variety of node architectures, operating systems, and even communication protocols was felt not to be an accident or a mistake but a permanent fact of life that should be supported up to a point. It might be argued that at least the protocols could be standardized into one set, but upon examining the list of application classes, it seemed to us that evolution and variety are required here too.

## 3.3.2. Resources/Technology Available

The position of the common carriers is rapidly changing, particularly in North America. We constantly read about deregulation, about ISDN, and about the coming voice/data integration. What is perhaps less visible, but for the present discussion much more important, is the astonishing rate at which buried optical fiber is being installed. It is also the case that this is a very economical technology. If one plots communication costs (e.g., cents/bit/second/kilometer) as a function of time over the last few decades, one observes a steady slow drop that has in the last few years become a

steep drop as fibers begin to be widely introduced. There is no corresponding revolution in switching technology to match that in transmission technology, but even here things are happening: slow-switching (digital cross-connects or "DACS" units) equipment at T1 (1.5 - 2 Mbit/s) and recently at T3 (45 Mbit/s) rates are being rapidly installed. In addition, switching at the raw fiber bandwidth of 0.5 to 1.7 Gbit/s is on the horizon.

Even though fibers are being installed first in intercity trunks and in local area networks, we feel that by the year 2000 fiber access between local premises and the central office (the "last mile" problem) should be widely available.

### 3.3.2.1. Bandwidth

At least two dozen important communication enterprises are installing buried single mode optical fiber, and over the next few years ten of them will have significant national coverage. Because of the observed magnetic attraction that buried fibers seem to have for backhoes, the more ambitious buried fiber carrier networks are 2-connected (and probably eventually 3-connected) between switching centers. Each installed fiber, although now being used at a single wavelength to send digital bitstreams at synchronous rates of 1.7 Gbit/s maximum, ought to be able to support 10 Gbit/s by the year 2000. The fibers are today usually installed in bundles, typically  $12 \times 12 = 144$  fibers in a bundle. One might get a capacity estimate for the year 2000 by noting that these numbers add up to 1 Tbit/s per bundle, and that when spatial redundancy and the multiplicity of carriers is included, most large cities should in the year 2000 be interconnected by terabits of raw digital capacity.

There are two wild cards to this prediction, whose realization could increase the capacity greatly even beyond these figures. The first is Wavelength Division Multiplexing (WDM). The intrinsic passband of the already installed single mode fibers is over  $10^5$  times the 10 Gbit/s that has been quoted above as the bandwidth of each fiber. In order to use more than today's one wavelength at a time, further advances in multispectral lasers and receivers (typically based on diffraction gratings) will be required. Operational WDM systems today use 2 wavelengths; figures up to 12 wavelengths characterize today's laboratory experiments. By the year 2000, the number of wavelengths simultaneously usable over fiber that was installed in the 1980s should be somewhere between 100 and the upper limit of  $10^5$ .

The other wild card is coherent optics: the use of the familiar heterodyne form of transmission and reception but at optical frequencies, rather than incoherent on-off energy detection. This should at least increase the signal-to-noise ratio and thus reduce the repeater spacing and may have other capacity-increasing consequences.

### 3.3.2.2. Switching

Today's DACS (Digital Automatic Cross-Connect Systems) frames, which are controlled by Telco operators on request from subscribers, will, under "software defined network" and ISDN plans of the major carriers, become subscriber-actuatable over the next few years. These provide switchover times of the order of seconds at a bandwidth granularity of DS-0 (64 kbit/s) or DS-1 (T1; 1.5 Mbit/s). DS-3 DACS equipment (45 Mbit/s) is expected to begin shipment next year. DACS switching speeds of the order of milliseconds appear to be no problem. Switching of entire per-fiber bitstreams is not available yet, but electronic GaAs-based switching or pure photonic switching is progressing and should become available in  $100 \times 100$  configurations by the year 2000. Today's switchover times of on the order of a millisecond may also characterize the photonic switching of 2000. Larger switches can be built by the usual concatenation of smaller switching elements.

### 3.3.2.3. Role of Satellites

By the year 2000, we do not expect satellites to play a major role in research communication networks. Their major usage niches are expected to be video broadcast and program distribution and access to hard-to-reach places not served by fibers (e.g., quick-install emergency recovery high-capacity links).

### 3.3.2.4. Communication Software

In order to utilize the high transmission bit rates available, today's long software path lengths will have to be reduced, and most of the communication function will be executed in high-speed hardware. In a high-speed packet switching service, for example, error recovery might be placed at the route ends and repacketizing along the route eliminated to reduce intermediate node processing.

### 3.3.2.5. Local Processing

There is always the question whether the growth of local processing power may not be actually leading to a reduction rather than an increase of the communication bandwidth required. Our projection, for example, is that workstation MIPs will increase from today's 4 to 100 times that by 2000. However, we also predict that while this is happening the need to intercommunicate will go up, simply due to the nature of research as an enterprise that requires large and apparently increasing amounts of intercommunication with colleagues, with shared resources, and so forth. The conclusion is that we can make no meaningful prediction on whether the increase in local processing MIPs will have a significant effect, either positive or negative, on the communication capacity required.

## 3.3.3. Conclusions

- (1) The technology of fiber communication will produce by the year 2000 a better than 1:1 ratio of the number of communication channels of all speeds available to the number of users that have a need for these channels. (This contrasts with 1:100 to 1:1000,000 in today's networks.)
- (2) The traditional packet switching approach, whose main purpose is to save on bandwidth by line sharing, may not be the right approach for many of the requirements of 2000. An examination of the table shows that the majority of the bandwidth requirements may be best served by circuit switching.
- (3) If prediction (2) is wrong and packet switching is widely needed in the year 2000, very fast packet switching, now only in the early prototype stage, will be required.
- (4) No new breakthroughs are needed. The technology is here already--only the numbers will change. Switching turns out to be no exception. Most of the requirements can be satisfied by switching times no faster than millisecond rates.
- (5) There will always be requirements for packet switching types of services. These include very rapid time multiplexing of messages for different destinations (e.g., many copies of the same short document to different addressees). Packet switching provides speed matching and allows for better management of highly dynamic connections as well as supporting the need for server resources (such as databases) to respond to many requests nearly simultaneously.
- (6) The homogeneity/transparency requirement also argues for a clear-channel circuit switch approach. The optimum research network of the future could be thought of as a homogeneous network of inhomogeneous systems.



### 3.3.4. Caveat

The 13 years between now and 2000 is a long enough period, not only for the technical evolutions we have postulated to take place, but for new unanticipated forms of user needs to spring up that will consume these resources. Historically, predictions of a future glut of some particular resource have often proved invalid. For example, even though real computer memory has been getting cheaper, faster, and smaller at a great rate, people are still anxious to have virtual memory, even on personal computers.

We are not able to predict a growth of need to match the growth of capacity, but it is quite possible.

## 3.4. SUMMARY

### 3.4.1. Considerations to the Year 1992

It should be a national priority to provide a ubiquitous computer communication capability to interlink the entire U.S. academic community. This should include a minimal basic service level for all as well as the possibility of a higher level of service when needed to support a particular application. This national communications infrastructure can be based on existing agency and academic networking activities. To be successful, significant new funding will be required. NASA, DoD, DOE, NIH, NBS, and NSF should be permitted and encouraged to seek ways of cooperating in dealing with issues of importance to the success of the research internet. These include the following:

- (1) Sharing of resources such as pooling of long-haul transmission requirements and multiplexed use of common 45-Mbit/s long-haul channels.
- (2) Joint engineering efforts to support the modeling, measurement, and management of the research internet.
- (3) Joint R&D efforts to address the problems of routing, congestion control, robustness, and performance in the internet and efforts to couple the results into migration plans from the DoD to the ISO protocols.
- (4) Coordination of user services and information clearinghouse functions.
- (5) Coordination of the development and operation of host/mail Name Servers and routers.
- (6) Establishment of a pilot program to procure (develop?) 45-Mbit/s packet switches to be put in place in anticipation of post-1992 needs.

### 3.4.2. Considerations to the Year 2000

Research and planning should begin now so that research networks can be based on the emerging technologies described above as these technologies become widely available.

### 3.5. WORKING GROUP PARTICIPANTS

Lawrence H. Landweber, Chairman  
University of Wisconsin, Madison

Harvey Newman  
California Institute of Technology

Bob Braden  
University of Southern California/ISI

Dan Van Belleghem  
National Science Foundation

Vinton Cerf  
Corporation for National Research Initiatives

Henry Dardy  
Naval Research Laboratory

David Farber  
University of Delaware

James L. Green  
NASA/Goddard

Paul E. Green  
International Business Machines

Bill Jones  
NASA/Ames

Tony Lauck  
Digital Equipment Corporation

James Leighton  
Lawrence Livermore National Laboratory

Barry. Leiner  
NASA Ames Research Center

Richard Mandelbaum  
University of Rochester

Raul Mazumdar  
Columbia University

John Morrison  
Los Alamos National Laboratory

Jan Postel  
University of Southern California/ISI

## 4. FUTURE STANDARDS AND SERVICE REQUIREMENTS

---

**Dr. Richard desJardins**  
Computer Technology Associates

### 4.1. THE IMPORTANCE OF STANDARDS

Historically there has always been a natural tension within the scientific networking community over standardization. On the one hand, standards promise widespread availability of vendor-supported products to facilitate economical interconnection of diverse computer systems. On the other hand, those who demand new and innovative networking capabilities often must develop new "standard" procedures, and may view existing standards as slowing or impeding the development of better ways of interworking.

To better understand the role that standards can and should play in the development of future networking for the scientific community, it is worth reviewing briefly how the computer industry's view of standards has evolved over the last decade.

### 4.2. THE ECONOMICS OF STANDARDS

Communication between two entities is impossible without prior agreement regarding the physical parameters of the communication channel and the meaning of symbols interchanged over the channel. While agreement can be negotiated individually for every communication, it is far more economical if many parties rely on conventions or standards that are widely shared.

More precisely, consider  $N$  users who wish to communicate using compatible hardware and software. In the absence of an effective standard, these users could develop bilateral agreements or translation packages that enable them to achieve compatibility. The total number of such agreements needed would be  $N(N-1)/2$ . Implementing and maintaining so many packages would be very costly to users. In practice, few such translations exist, with the result that users are frequently unable to communicate altogether. The adoption of a standard replaces many separate bilateral implementations with a single implementation of the standard for each of the  $N$  systems. Thus, in the presence of a standard, implementation costs are of  $O(n)$  rather than  $O(N^2)$ .<sup>1</sup>

Users value standardization because it reduces their cost of establishing communication between heterogeneous systems. Not only are the costs of bilateral translation systems avoided, but the costs of managing and maintaining many diverse packages are avoided as well. In the absence of a standard, it is often the user who must pay for custom implementation of translation packages. In the presence of a standard, the vendor will typically make an investment in its implementation, allowing the user to buy commercial off-the-shelf equipment and software. Standards often lead to increased vendor competition based on price rather than compatibility, with attendant benefits to buyers.

Standards, by holding out the promise of a substantial market for an interface, can persuade semiconductor manufacturers to invest in low-cost VLSI implementations.<sup>2</sup> Standards can also reduce costs by allowing vendors to achieve economies of scale.

The benefit of being able to communicate easily with others increases the value of a computer system to a buyer. The larger the "network" of users with which a particular system can communicate, the greater is its value to the buyer. This increase in value leads in turn to increased demand for computer systems.<sup>3</sup>

Recognition that standards can lead to larger markets has led firms to support standards even when, as a result of standardization, there is intensified competition based on price. For this reason, the computer and communications industries worldwide have invested an estimated \$250M\* since 1977 in the development of a broad set of standards for Open [computer] Systems Interconnection (OSI) under the aegis of the International Organization for Standardization (ISO). It is likely that the Defense Advanced Research Projects Agency (DARPA) and the scientific networking community have invested an equivalent amount in their standardization program over the past 15 years; certainly IBM, for example, has invested substantially more, probably upwards of \$1B in its Systems Network Architecture (SNA) program over the past 15 years.

Notwithstanding the benefits of standards, achieving consensus is not an easy process. Various firms or buyers may attempt to impose de facto standards. When two or more technologies are competing to become the standard, those who make an early decision may suffer from orphan investments in an unsupported technology. In some cases, a superior technology will fail to be adopted because the costs of abandoning an existing technology appear to be too high, or because the absence of an existing network of users makes the new technology appear insufficiently attractive.<sup>3</sup> (Consider the problems of the Dvorak typewriter keyboard,<sup>4</sup> or the slow rate of conversion from bisync to SDLC.)

On the other hand, adoption of a standard can be accelerated by coordinated actions on the part of buyers (e.g., GM and the Manufacturing Automation Protocol group<sup>5</sup>), by "sponsors" of a new standard who lower the costs initially to persuade users to switch,<sup>6</sup> or by coordinating vendor statements of support for standards, which helps persuade each of them that they will not be alone in switching to a new technology.<sup>7</sup> All of these methods have been used successfully in the network research community.

In summary, standards can provide significant economic and performance benefits to both users and vendors. At the same time, the shift to standard solutions can entail costs to early adopters of nonstandard technology and is difficult to coordinate.

### **4.3. HISTORY AND TRENDS IN NETWORK SERVICES AND STANDARDS**

In this section, we take a brief look at the history and trends in computer network services and standards.

#### **4.3.1. History of Computer Networking Standards**

The concept of packet switching is generally credited to Paul Baran in a report for the Rand Corporation in 1964.<sup>8</sup> Baran did not use the term "packet switching," however. The credit for that term goes to Donald Davies who by 1965 was at the U.K. National Physical Laboratory working on a single node packet switch. Davies never developed a multinode switch, however, and so did not have to face the tandem store and forward problems inherent in modern packet switched networks.

---

\*This figure includes all costs of attending meetings plus activities conducted between meetings within individual firms in support of standards. Investments in implementations of the standards would add substantially to this figure.

DARPA in the U.S. carried out a two node host-to-host packet switching experiment in 1966 between the Q32 at Systems Development Corp. (SDC) in Santa Monica, California, and the TX2 at MIT Lincoln Laboratories in Massachusetts. Lessons learned from this experiment formed the basis for the Arpanet Request for Proposal (RFP) in 1968, which led to the delivery by Bolt, Beranek and Newman (BBN) of the first ARPANET packet switching node to UCLA in 1969.

By 1972, two commercial packet switched networks--Tymnet in the U.S. and RETD in Spain--were in operation, along with the private international airline reservation network (SITA). By 1974, research or experimental packet networks were in place in France (RCP and Cyclades) and Japan. RCP was the predecessor of the French commercial packet switching network Transpac. Cyclades--about 6 months behind RCP--was spearheaded by Louis Pouzin at IRIA (later INRIA). A key member of Pouzin's team at IRIA was Hubert Zimmermann, who in 1978 became the chief architect of the OSI Reference Model.

By the mid 1970s, public packet switched networks were being planned in the U.K. (EPSS), France (Transpac), the U.S. (Telenet), and Canada (Datapac). Two pan-European projects--Euronet and the European Informatics Network (EIN)--were also on the drawing boards.<sup>9</sup>

Each of these experimental networks used a different set of standards or conventions for the connection of hosts to the network. Each such standard required that appropriate software be written for every brand of host computer that would eventually be attached. The cost of writing such software for many different networks would be enormous.

Two approaches to dealing with the proliferation of networks were adopted by various groups. The builders of public packet networks sought to define a single standard interface for the connection of a host computer to a packet network. This standard would serve the purpose of inducing computer manufacturers to develop the software that would support the emerging public packet data networks. Without such support from the computer vendors, the packet network vendors would have to provide the software to make their networks usable. Clearly this would be an expensive and difficult proposition. Defining such a standard interface between a carrier provided network and customer equipment was typical of the way the telecommunications community operated. Accordingly, the CCITT appointed a Special Rapporteur for Packet Switching Standardization in July 1974.

The international scientific community approached the problem differently. Rather than standardize each packet network, researchers in the DARPA community focused on standards for end-to-end communication across several concatenated networks.<sup>10</sup> A simple Internet Protocol (IP), which required minimal datagram service from each individual packet network, would underly an end-to-end protocol (Transmission Control Protocol--TCP), which would provide high-quality service from host to host.

In designing a network interface, a key choice is the division of responsibility for reliability and packet sequencing between the network vis a vis the attached computers. In the DARPA internetworked community environment, where many heterogeneous networks would be involved, it was infeasible to require all existing networks to change their interfaces and internal designs. Moreover, some of the network technologies were inherently lossy (e.g., Ethernets, packet radio, packet satellite).

Therefore the DARPA researchers concluded that in the heterogeneous internetworking environment it was technically necessary to assign the responsibility for end-to-end reliability to the hosts and thus to minimize the required network functionality by specifying only a network datagram service (which became the basis for the IP). Working through the International Federation for Information Processing

(IFIP), representatives of the research community in the U.S., the U.K., and France developed a common proposal for an end-to-end transport protocol.<sup>11</sup>

In contrast, the public network operators desired for commercial reasons to maximize the value added by their networks. A vertical circuit interface would put responsibility for packet sequencing and reliability in their networks. Moreover, it was not clear how such commercial requirements as varying grades of service, "collect calls," and closed user groups could be provided to customers through a datagram interface. Finally, a virtual circuit interface would be easier to explain and market to customers accustomed to point-to-point private lines.

Because the CCITT works in 4-year cycles, it was necessary for the public network vendors to reach an agreement by September 1975 if a standard was to be ratified during the 1976 plenary. Thus, by the time IFIP submitted the compromise proposal from the research community, it would have been too late even if the CCITT community had been inclined to adopt a datagram-oriented approach. Thus this single opportunity to bring the two approaches together never materialized.

DARPA never made it a priority to participate in the national or international standardization efforts because its primary charter was to meet military needs. DARPA continued to refine the TCP/IP standards, which were formalized by 1978 and were adopted as DoD procurement standards in 1982. DARPA supported many of the academic research institutions to participate in the internet project, which led to the interconnection of all the research networks in the U.S. around TCP/IP as a basis. DARPA made TCP/IP software available in the public domain through the Berkeley UNIX distribution mechanism.

The military Defense Data Network (DDN) also made use of TCP/IP, although many agencies waited until DDN arrived before pushing hard to obtain TCP/IP protocol implementations commercially. DoD actually sponsored some of the commercial development for use by the military.

In any event, both the research and the commercial computer community recognized that the network standards--whether X.25 or TCP/IP--fulfilled only a part of the need for computer communication standards. Standards were needed for process-to-process communication, not just host to host. The most important of these application protocol standards were for file transfer, electronic messaging, and virtual terminal. DARPA and the research community had developed just such a trio: File Transfer Protocol (FTP), Simple Mail Transfer Protocol (SMTP), and TELNET for network virtual terminal and login.

In the commercial standardization community, ISO began the development of a suite of international standards for OSI with an initial meeting in March 1978. OSI became a truly universal standardization program when CCITT officially approved OSI as a collaborative program with ISO in 1980.

The OSI activity began with an unusual step: the development of a layered reference model that described an architecture within which the various existing standards such as X.25 and the proposed standards such as file transfer would fit. The development of an architecture permitted multiple working groups to proceed in parallel to develop standards at the various layers.

The OSI effort struggled with the same problem of virtual circuits versus datagrams that had been fought over in the CCITT X.25 effort several years earlier. This time, however, a compromise solution was adopted in which both datagram and virtual circuit approaches were supported.

The principal U.S. government funded participation in OSI was not from the research community but from the National Bureau of Standards (NBS) and the National Communications System (NCS). Not only did these agencies send their own employees, but NBS contracted with ARPANET developer Bolt Beranek and Newman (BBN) to prepare materials for consideration by ISO working groups. Other agencies such as NASA, Navy, and Army provided support for individual employees to participate.

One might ask why TCP/IP was not presented as the solution to the recognized need for supporting a connectionless network (i.e., datagram) service and a connection-oriented transport protocol within OSI. In fact, TCP/IP was presented to CCITT as a possible international standard but was rejected largely because of its U.S. military origin. In discussions within ISO concerning the possibility of using TCP/IP, many non-U.S. manufacturers were reluctant to embrace in toto any protocol for which U.S. firms had a significant lead in implementation as a result of U.S. government funding.

Moreover, by the time detailed discussions got under way in 1981-82 on an ISO-CCITT transport protocol class TP4 similar to TCP, the state of the art had advanced somewhat. Thus, for example, TP4 was designed to allow transport acknowledgments to be returned on any connection going back to the originating host, not just on the particular connection being acknowledged. Also, it was a requirement that TP4 share transport connection establishment and encoding methods with the other ISO-CCITT transport protocol classes. While these influences caused the details of TCP and TP4 to differ considerably, they were functionally equivalent because TCP was an ancestor and prototype for TP4. In fact, many of the technical experts who developed TP4 and other OSI protocols within ISO had considerable experience with TCP/IP because they were the early implementors of TCP/IP.

The TP4 standard was approved by both ISO and CCITT in 1984. The emergence of this international standard with function and performance nearly identical to TCP forced the DoD to confront again the issue of standardization of its own protocol family. The DoD goal has always been to buy compatible networking software off the shelf from commercial suppliers. By supporting the TCP/IP protocol suite as a de facto standard, DoD hoped to achieve that goal. In fact, by 1987, a substantial TCP/IP support industry has developed. Most computer manufacturers now support third parties.

On the other hand, OSI was adopted by ANSI in 1980 as the basis of American National Standards for computer networking protocols. By 1984, the major computer manufacturers such as DEC and Honeywell were announcing that they intended to support OSI as the basis of their networking architectures. Furthermore, both the European community and Japan adopted OSI as their architecture standard. By 1984, the momentum that was gathering behind the ISO architecture and protocols worldwide began to raise questions as to the long-run costs and benefits to the DoD of staying with TCP/IP versus migrating to OSI.

At the request of the DoD and NBS, a National Research Council panel reviewed the issue in 1984-85 and recommended that the DoD move rapidly to put TP4 on an equal footing with TCP. In March 1985, the DoD announced that it intended to study the issue further before proceeding. During the following 2 years, the DoD and NBS jointly developed and are now implementing a transition plan aimed at moving not just to TP4 but to the entire ISO protocol stack from the ISO internet sublayer on up. Beginning in 1988, the DoD plans to authorize its agencies to procure ISO-compatible systems as well as TCP/IP-based systems. Beginning in 1990, compatibility with the ISO protocol suite is planned to be mandatory.

Among civilian agencies, a Government OSI Procurement (GOSIP) policy is being developed that would require all government computer network procurements in the future (including those of the DoD) to be capable of using the OSI protocols. TCP/IP support is also authorized in the interest of interworking with the internet community.

#### 4.3.2. Trends in Standards Development

During the nearly 20 years of computer networking history recounted briefly above, the process of information technology standards development has changed radically. Currently, the development and deployment of standards can be seen as encompassing six distinct phases.

The first phase is the development of an overall Reference Model or architectural framework to guide the development of a related set of standards. Beginning with X.25, standards bodies began to use reference models to divide the protocol standardization problem into layers of functionality. Layering enables separate groups to work on the development of related standards in parallel.

In the second phase, terms and parameters are defined carefully and agreement is reached on the functions to be performed by the standardized product or procedure. This agreement is set down in a standard Service Definition.

In the third phase, the standards body rigorously describes in a Protocol Specification the specific methods used to accomplish each communication function. The Protocol Specification always contains a conformance section describing precisely what must be implemented for a vendor to claim adherence to the specification. Test methods may also be specified. These specifications are written entirely or in part in a formal description language.<sup>12</sup>

Increasingly, standards agreements include numerous options to satisfy different participants. Consequently, two products, both of which conform to the same standard, may not interwork at all or may interwork poorly if the vendors have implemented different options.<sup>13, 14,\*</sup> Thus in the fourth phase, user or government sponsored groups specify precise parameter values, classes, and options as a way of further reducing the diversity.<sup>15</sup> These procurement specifications are called functional standards or profiles. The most important of these forums in the U.S. is the NBS OSI Implementors Workshop, which has been operating for several years and has produced agreements among vendors and users as to which protocol options and guidelines to adopt to ensure interoperability.

In the fifth phase, firms develop products conforming to the standard and bring the products to market. Each firm attempts to position itself in the marketplace in a manner that supports the standard but adds differentiating value such as cost, performance, flexibility, or user friendliness.

Given the complexity of most information technology standards, however, a sixth phase, conformance testing--by vendors, users, third parties, or the government--is necessary to ensure that the standard has been fully and correctly implemented.<sup>16</sup> Only then can buyers be confident that a purported standardized product will operate as expected. In the U.S., the Corporation for Open Systems (COS) has been established to develop conformance testing procedures for OSI products. Similar organizations have been formed in Europe and in Japan.

---

\*For example, the X.25 standard allows several different packet lengths. The British PSS uses 64 and 128 octets, while AT&T supports 128 and 256.



Because of the economic importance of standards, firms have been reluctant in many cases to bring out new products until relevant standards are established, whether international voluntary standards or de facto industry or user standards. Thus, after its experience with Beta versus VHS, Sony delayed introduction of compact audio disks until all the major firms had agreed to adopt one standard. Unfortunately for Sony, this early standardization enabled low-cost producers to quickly and brutally undercut Sony in the marketplace. Nevertheless, for users, the result is that compact disk players introduced by Sony at \$1000 sold just 2 years later for \$300 and now have dropped even lower.

In the field of information technology, especially in computer communications, standards are becoming much more anticipatory and are being supported more openly by coalitions of vendors, users, and governments. Firms are generally attempting to align their product developments with these emerging standards as rapidly as possible to minimize the costs of later conversion. Users benefit from this situation, which will continue as long as the firms individually perceive the risk/reward ratio to be favorable. In the meantime, third parties will appear in the marketplace offering products to meet any unfulfilled demand for standardized solutions.

#### 4.4. LESSONS LEARNED IN NETWORKING STANDARDIZATION

In preparation for a major U.S. research networking initiative, it is useful to review some of the lessons learned in networking standardization over the past decade and a half, especially concerning the relationship between theory and practice in networking services and protocols.

There is a major difference between the requirements of research and the requirements of the commercial sector. Networking research is by definition the leading edge in the development of new networking concepts and technologies. Networking research is concerned with solving a technical problem, not with embedding that technical solution in a product. In some cases, this requires building a prototype to adequately investigate the problem, but that prototype is seldom sufficient for a product that has to be supported in the marketplace.

For example, one might build a "toy" database system to investigate problems of distributed database systems, but one would not go to the trouble of building a complete DBMS to investigate the problem. The goal is different. The goal of the research would be to investigate the problems of distributed databases (with as few independent variables as possible), not to build a working product. Once the problems have good solutions, a product can be considered. One of the biggest problems in these days of very expensive research is how to lessen the time and cost of the transition from research to product.

The history of network service and protocol standardization is a very good object lesson in how to move networking research results to standards on a global scale. Here are the ten major lessons that have been learned.

**First, don't codify current practice.** Take the latest research results, add the best engineering judgment available, and produce standards aimed at a point in the future. Given the rate of technological change, the complexity of the required standards and the time necessary to get international agreement, any other approach would produce a document of only historic interest that would either be ignored or would act as a drag on technology development.

**Second, take a systems approach.** Define an architecture for the area to be standardized, so that an entire set of standards can be designed to work together. This approach ensures that the final standards will be compatible and makes it possible to determine clearly what does and what does not

need to be standardized. (One of the most common mistakes in developing standards is to confuse what belongs in a standard with what belongs in an implementation design.) The systems approach allows parallel development of standards to speed the process and also provides the framework for managing the development.

**Third, take an incremental approach.** Develop the architecture and the standards so that they can be extended. Don't try to provide the architecture for everything before starting to work on any of the standards. There are two reasons for this: it makes useful standards available sooner, and it provides feedback from the early, more detailed work to ensure that the architecture remains close to what is pragmatic.

**Fourth, define standards with sufficient rigor** that conforming implementations can be developed without considerable interaction with people who already have implemented the standard. If the only way to get a conforming implementation is to work with someone who already has one, this will be a major impediment to the adoption of the standards. A rigorous specification also provides a basis for rigorous analysis and validation of the standard, and a firm reference for verification and conformance testing of implementations. OSI has pushed the development of Formal Description Techniques. Once again software has lagged behind hardware. No one would consider designing VLSI by hand anymore, but we are still designing complex communications protocols by hand.

**Fifth, avoid isolating research.** Networking research must be kept close to its natural soil: real networks. There are two principal reasons for this: to understand what sorts of networks are required, and to ensure that the latest research results are incorporated into real networks. Otherwise the research is useless. Research must be done with the technology transfer phase (to product) in mind.

**Sixth, have a plan for networking research.** No single research problem is isolated from all others, especially in networking. A clear understanding of the relationships among the problems and how some solutions can be built on top of other solutions (or affect other solutions) can do much to improve the effectiveness of the research. Doing research according to an architectural framework provides a powerful focus on the state of the art and on the problems of transferring the technology.

**Seventh, do not overly constrain implementation.** This allows implementations to be tailored to particular research environments, while promoting further research within the framework.

**Eighth, recognize that the standardization process is a people process.** The people involved must be very up to date in the field, understand the requirements in sufficient depth to recognize a solution that takes into account all the requirements, understand the economic and political issues important to each participant, not just be an observer but help formulate win-win solutions, and have a very good sense of the marketplace in which the standards must flourish. These talents, collectively known as electro-political engineering, are crucial to success in international standardization.

**Ninth, distinguish between research networking and networking research.** Research networking is networking for the scientific community and provides a service to a wide variety of researchers in many fields. The requirements of most of these researchers is for a highly reliable basic networking service. Thus most researchers want the network to be as invisible as possible. In contrast, networking research is research in the field of networking. It is aimed at gathering a better understanding of the principles governing the behavior of networks and pushing the development of network technology in promising directions. Networking research includes both lower layers research, aimed at the development and understanding of new network technologies, and upper layers research, aimed at the development and understanding of the principles and techniques of distributed systems.

**Tenth, distinguish between products and prototypes.** It is natural that by the time research is to be incorporated into a commercial product (standard) more will have been learned and other issues will have to be taken into account so that it is very likely that the end result will not be precisely the same as the research solution. Thus it is important that networking researchers ensure that their results are transferred to the production networking environment. Since it is likely that the final standard will be different from the research prototype, the impact on research systems can be minimized if the research is done in the context of a standard architecture, so that experimental protocols can be swapped out and replaced by standard protocols as needed, or in a few cases, so that the standard architecture itself can be adjusted in the light of experience and emerging technologies. For example, in the office document area, current NSF research on compound documents involving text, equations, spreadsheets, figures, and images is building on top of existing ISO and CCITT standards rather than starting from scratch or building on a proprietary solution.

#### **4.5. NETWORKING SERVICES FOR THE SCIENTIFIC COMMUNITY**

From a historical perspective, the ARPANET network services were the basis for end system to end system (host computer to host computer) data communications for the scientific research community. ARPANET has now been interconnected with other systems, which provides communication services to a large number of universities, federally contracted research and development organizations, DoD establishments, and other organizations that extend across the U.S. and Canada as well as Europe. These networks connect heterogeneous end systems that use standardized services, with the goal of making the heterogeneity transparent to the user.

We emphasize again the distinction between research networking and networking research, i.e., between providing a stable base for researchers who use networks to support research in other fields, vis a vis using the network itself as a research vehicle. The nature of research oriented toward advanced networking capabilities characteristically brings with it some instability and perhaps unreliability in communications services. In this discussion we are interested in identifying standard networking services that should be offered as the stable and reliable basis of research networking.

The actual services provided by present networks are most often limited to electronic mail or messaging, file transfer, and remote login or virtual terminal services. In practice, these services depend on either a reliable end to end data transport service or a reliable network connection service.

The DARPA data communications protocols that provide the services of electronic mail, file transfer, virtual terminal/remote login and reliable data transport are used ARPANET as well as in many other networks in the internet system. However, the trend in commercial networking is toward international standards, more specifically the OSI standards. Therefore, it is necessary to consider the strategic and practical issues involved in identifying networking standards for future networking services for the research community beyond about 1990. The remainder of this report provides our recommendations for the needed strategy.

#### **4.6. STANDARD NETWORKING SERVICES**

The need for a stable base of computer networking communication services for the scientific research community must be recognized as being of prime importance. The argument for stability is simply to enhance the productivity of researchers rather than burden them. The stable base should start with what is now in place, that of electronic mail/messaging, file transfer, and remote login/virtual terminal service, and the underlying reliable data transport/connection service, based on the DARPA standards. (Note: This does not address the question of whether those services are now offered with sufficient

capacity for the scientific research community. Questions of capacity are discussed elsewhere in this report.)

Currently, computer subnetworks are based on either telephone lines or local area networks (LANs) as the underlying physical facilities. Standard local area networks provide data rates up to 10 Mbit/s while subnetworks based on telephone lines currently provide much lower rates, typically 56 Kbit/s and below. Higher speed services are the subject of intensive activity, both in the provision of high-speed digital services today and in planning for the standard services of the future.

Standards for the Integrated Services Digital Networks (ISDN) of the 1990s are being developed by CCITT; these feature subscriber basic rates totaling 144 kbit/s on three channels through a single mini-connector similar to today's modular telephone jack, as well as providing primary rates of 1.544 Mbit/s and above. Standards for the Fiber Distributed Data Interface (FDDI) are being developed by IEEE; this features a basic data rate of 100 Mbit/s. These new types of subnetwork facilities will be needed to adequately support supercomputer networking at reasonable cost. These facilities will also be needed to exploit the rapidly growing capabilities in long-haul fiber optics and satellite systems. Some additional work will be required in the standards areas to tune the performance of both DARPA and OSI protocols to these emerging high performance underlying facilities, and to develop new protocols that fully exploit these new capabilities.

In the near future, some additional services beyond the basic ones mentioned must be developed to support the research community. The first of these is the ability to exchange documents between end systems via network connections. In this context a document is any grouping of data that is meant for human consumption in its primary representation. Text in different fonts and different levels of processing, images both raster and graphic based, scientific equations, spreadsheets, voice representations, and animated graphics sequences are some examples of these groupings of data. Present abilities for such exchange are generally limited to end systems that have pre-arranged understandings for the content of the documents and a corresponding ability for processing the contents. The diversity of internal representations of data as well as the proliferation of editors, word processors, computer aided design (CAD) representations, etc., associated with the massive numbers of different types of computers demand an easy method for the intelligent exchange of documents.

The OSI architecture and its associated standards give some ability for simple to moderately complicated message exchange but the needs of the research community extend to the extremely complex as already noted. High-priority research efforts should be directed at advancing document description technology such as the Office Document Architecture (ODA) and "data description language," with the goal of integrating the document descriptions into the communications architecture either in the end-systems or as a network service. This effort should explicitly avoid representations.

The second category of services needed in the immediate future are known as directory services. The analogy of the "white pages" telephone book to look up addresses of users of networks is a minimal example. Computer networks associated with the present research community have tens of thousands of users. To allow one user to communicate with another, the target user address must be known or a means for lookup must be provided. The capability for the address identification for a "user object" exists now in specific implementations. However, the capability lacks extensibility across heterogeneous end-system implementations due to the lack of standardization.

Given that the nature of networks resembles an information repository as well as a telephone system, users are not the only objects of interest. Scientific literature or documents are stored in multiple end systems just as they are stored on different shelves of a library. An analogous capability to a card catalog for locating the storage location of such document objects is necessary to facilitate the

scientific community's use of networks to accomplish scientific research. As an extension to this, various attributes associated with objects need to be accessible. An example might be the need to locate a computer simulation model for virus growth patterns that resemble the attributes of AIDs or a multiple precision mathematical routine for calculating satellite orbital trajectory. The analogy of "yellow pages" is also a needed service associated with objects significant in networking. Often the attributes of an object are known but the specific name is not.

Directory services are also important to the operation of the network itself. As an example, the management function of controlling access to network resources or objects may need to make use of subdirectories of authorized users.

The lack of standardized directory services in the research networking community in each of the application areas discussed demands a high-priority effort in the near term. ISO and CCITT are developing a collaborative directory standard that would apply to these applications.

Two additional areas are just beginning to become mature enough for inclusion into a standard service base. First, the integration of **voice and video conferencing services** into networks for the scientific community should be studied as the capacity of networks are expanded. Opinions differ on whether the OSI architecture is appropriate for voice and video, but in any event, either the OSI architecture should be used or more appropriate architectures should be developed. In no case should standardization of voice and video be done without an architecture, and whatever architecture is adopted should take account of international standardization activities in these areas.

Second, the general topic of **distributed processing services** should be considered. Of immediate importance is the need to ensure that directory services associated with distributed processing objects are included in the requirements for the networking directory architecture. Distributed processing associated with networking allows a particular end system to make use of the processing capability of other end systems attached to the network as if they were an extension of itself. In the ideal case, this service virtualizes the network and participating end systems into a single entity and delivers the service transparently to the using entity, be it an actual user or otherwise. The objects associated with this virtualization have multiple attributes, locations, etc. and are of concern in a directory services sense. Research and standardization efforts in distributed processing for both the directory services application as well as the larger general case are becoming critical. Again, distributed processing service standards in the research networking community should take account of standards such as commitment, concurrency, and recovery (CC&R), remote database access, transaction-oriented protocol, and common procedure call that are under development within ISO.

Finally, many application services such as tools, databases, expert systems, and supercomputer support may quickly evolve to become candidates for standardization in the near future.

#### **4.7. SERVICE RELIABILITY**

Given that stability of network services for the scientific research community is a requirement, the reliable delivery of the services must also be a requirement. Apart from the reliability of the end system and the network hardware and software components themselves, several mechanisms are precursor to providing reliable networking service.

Within the networking standards community associated with OSI (and to a large extent the present Arpanet/DoD standards community), the protocol standard specifications are accompanied by **standard service definitions**. These standard service definitions provide the consistency and uniformity

specifications that allow implementations for heterogeneous end systems as well as for network service components to provide reliable services. The completeness and quality of these service definitions are therefore crucial, and the research community must contribute its expertise to help assure this.

A parameter or attribute of most standard service definitions is **quality of service (QOS)**. A familiar example is that of advertised bandwidth and distortion (and associated sound quality) associated with a telephone connection. These quantities can be measured, and they constitute attributes by which a consumer can judge the performance of the service provider. Generally, for example, advertised telephone bandwidth is in fact delivered and distortion is within specified limits.

A network service provider should advertise QOS parameters such as bandwidth, transport delay, and error rates. The variability and outright deficiencies in QOS assurance are the cause of much of the instability of current network services. A significant research effort is appropriate to establish methods for ensuring a uniform quality of service.

Another important reliability characteristic of networks as well as of the heterogeneous end systems attached to them is a **high degree of autonomy** in their operation. The administrative procedures applied to the operation of end systems and networks are themselves heterogeneous. This complicates issues such as uniform quality of service if the autonomy allows adverse impact to other users caused by improper administration. Management and administration of the large networks formed by interconnection of multiple smaller networks typical of the research community are a significant problem and should be carefully studied for improvement.

Within the area of network management, additional services are required to facilitate the research community's efforts. Concerns for **privacy or confidentiality** of certain information are easily understood. Privacy can be provided interior to the network (as contrasted with the more difficult case of end to end) with current technology. The security issues within end systems are generally not network service provider problems and must be dealt with by local systems owners and operators and by the networking community as a whole. Efforts are under way to define the needed **end-to-end security services** within the research community and they should be encouraged to produce results as soon as possible.

Other security services such as user authentication (proof that a user is in fact who he or she says they are) and non-repudiation of originator (sometimes called digital signature service) are least partially in the realm of network provider services. They are not available now and need immediate efforts to establish them. Again, ISO and CCITT are active in these areas, and the efforts of the research networking community should take account of the international standardization activities.

Finally, **accounting and charging services** should be provided, not so much to discourage user subsidies (which are frequently justifiable as a research facilitator) as to promote efficient management practices on the part of network service providers.

It is evident that a number of the networking services that should be provided to the scientific research community are not yet available. The starting basis for a number of these new standard services are available such as the Office Document Architecture (ODA) efforts mentioned above. In order to evolve the basis of these services, aggressive approaches must be undertaken. The **funding of prototype implementations** of preliminary services with the express purpose of evolving such services to a production capability has had success in the past. If results of prototyping are fed from the research and development community into the standards community, a synergy benefit could be established. If the prototype implementations are developed with the user community in mind, a market base can be established as well.

An example of such a development and demonstration process can be found in the NSF Express project. This effort will provide an online research proposal submission capability that allows the interchange of documents containing complex text including equations and graphics between end systems running different editor/word processors. The basis for this capability is the integration of ODA into each of the editors. At a minimum, projects to develop or evolve the services discussed in each of the above paragraphs, such as document interchange, directory services, management and administration services, etc., which are not yet available, should be undertaken immediately.

#### 4.8. TRANSITION STRATEGIES

As discussed in the historical perspective on network services and standards, DoD has announced that it will commence a transition from the DARPA protocols and services to OSI in the immediate future. Given that DoD-supported networks comprise a large portion of the support base for the scientific research community, **strategically planned transition for the entire scientific research community to OSI standards including ISDN should commence immediately.** This strategic plan should outline the timing for such a transition as a priority so as to make clear the requirements to the full community. This recommendation is supported in the report of the National Research Council on "Transport Protocols for Department for Defense Data Networks." Early transition planning can result in major savings if it is done prudently. The transition planning effort for the scientific community must be closely coordinated with the efforts already under way in DoD to minimize duplication of effort and maximize mutual benefit.

In addition to the planning, **real specific transition mechanisms should be actively pursued.** DoD has identified three interoperability mechanisms to allow communication between DoD and OSI end systems during the transition period. First is the ability to **support either the DoD or the OSI reliable transport service.** This requires that the network intermediate system components provide a "dual IP" capability, i.e., be able to recognize both the DoD Internetworking Protocol and the OSI Internetworking Protocol and to apply routing procedures specific to the protocol type received. An alternative approach is to encapsulate DoD Internet Protocol data units internal to OSI Internet Protocol data units (or the reverse). The second transition mechanism is the use of **translating application gateways.** Such devices would map the electronic mail/messaging and file transfer applications to their counterparts in the opposite suite of protocols. This will allow the basis for interoperability among the basic services discussed above. The third transition mechanism to be used by DoD will be the use of **end systems that can operate either suite of protocol** on demand from the user. Providing for this type of implementation in a large number of the various types of end systems may be prohibitively expensive as a universal solution but will be achievable at moderate cost in a great many cases by the use of public domain and commercially available software solutions. These transition mechanisms being employed by DoD are equally applicable to the research networking community, and their development and use should be supported.

A fourth transition method developed by Northrop Corporation uses a **thin sublayer on top of TCP to transform it into the OSI Transport Service** as provided by TP4. With this approach, upper layer OSI protocol stacks can ride on top of existing TCP/IP networks alongside existing DoD application protocols.

A fifth method to facilitate transition in the research community is through **funding of public domain implementations of the OSI protocols.** As an example, if such implementations were developed using as a basis the "open architecture" operating systems popular in the academic community such as Berkeley UNIX, the low cost of the capability would provide easy entrance to the newer services and therefore would accelerate transition and promote networking research based on the new standards.

There will almost certainly be some useful transition mechanisms that would not be commercially viable as products. The planning done for transition should identify such needed software and hardware, and high-priority efforts should be put in place to provide incentives or full subsidization as necessary.

Within the activities associated with transition, as well as for research on advanced services, some **guidelines for future methods of transition** are also necessary. DoD has suggested that the current application protocols that provide electronic mail/messaging, file transfer, and virtual terminal services should not be modified individually to run over OSI based transport services as this would induce additional transition steps. This guideline should be considered for the research community. On the other hand, running existing FTP or TELNET software over an OSI application association would be a very simple matter and might offer in particular cases a useful ad hoc services for the research community suggests that applied research on advanced communications services should **build on existing and emerging standards where possible** and should avoid the use of production networks if an adverse impact on basic communications services is likely. Networking research areas aimed specifically at improving research networks and affecting future standards should be encouraged.

Finally, **American networking technology should be promoted as the basis of international standards.** The alternative is to accept European and Japanese leadership if our networking research knowledge is not being made available due to lack of participation. If European and Japanese researchers and manufacturers define the technological basis of international standards, the result will be a world networking environment in which American manufacturers are not able to fully exploit their technological excellence.

#### 4.9. WORKING GROUP PARTICIPANTS

Richard desJardins, Chairman  
Computer Technology Associates

Shirley Radack  
National Bureau of Standards

Roger Cheung  
Hewlett Packard

Marvin Sirbu  
Carnegie Mellon University

John Day  
Codex/Motorola

George Sullivan  
Defense Communications Engineering Center

Debra Deutsch  
Bolt Beranek and Newman Inc. Laboratories

Additional contributors to historical summary:

Ronald Ham  
Digital Equipment Corp.

Vinton Cerf  
NRI

Jonathan Katz  
The Analytic Sciences Corp.

Barry Leiner  
RIACS

Dan Lynch  
Advanced Computing Environments

Richard Pietravallo  
Digital Equipment Corp.



#### 4.10. REFERENCES

1. J. McQuillan and V. Cerf, eds., "A Practical View of Computer Communications Protocols," IEEE (1978).
2. M. Sirbu and K. Hughes, "Standardization of Local Area Networks," Proceedings of the 14th Telecommunications Policy Research Conference, Airlie, Virginia, April 1986.
3. J. Farrell and G. Saloner, "Installed Base and Compatibility: Innovation, Product Preannouncements, and Predation," *American Economic Review* (December 1986).
4. P. David, "Understanding the Economics of QWERTY: The Necessity of History," in *Economic History and the Modern Economist*, W. Parker, ed. (Basil Blackwell, 1986).
5. B. Schatz, "GM, NBS Push ISO Plan," *Datamation* 9, 66 ff (June 15, 1984).
6. M. Katz and C. Shapiro, "Technology Adoption in the Presence of Network Externalities," *Journal of Political Economy* 94, 4, pp. 822-841 (August 1986).
7. "Major Companies Endorse European LAN Standard," *Control Engineering* (August 1982).
8. P. Baran, "On Distributed Communication," Rand Corporation Technical Report RM-3420-PR (August 1964).
9. M. Sirbu and L. Zwimpfer, "Computer Communications Standards: The Case of X.25," *IEEE Communications* (March 1985).
10. V. Cerf and R. Kahn, "A Protocol for Packet Network Intercommunication," *IEEE Transactions on Communications* (May 1974).
11. V. Cerf, A. McKenzie, R. Scantlebury, and H. Zimmerman, "Proposal for an Internetwork End-to-end Protocol", IFIP, 1975
12. G. Bochman and C. Sunshine, "A Survey of Formal Methods," in *Computer Network Architectures and Protocols*, P. Green, ed. (Plenum, New York, 1982).
13. M. Hess, M. Brethes, and A. Saito, "A Comparison of Four X.25 Public Network Interfaces," IEEE ICC '79 Conference Record, pp. 38.6.1-8 (1979).
14. "Analysis of European X.25 Networks," European Computer Manufacturers Association Technical Report 15 (September 1983).
15. "Intercept Recommendations for the OSI Transport Layer," U. K. Department of Trade and Industry, Information Technology Standards Unit Technical Report TG 102/1 (March 1984).
16. Steve Pollili, "Independent Standards Testing for MAP," *MIS Week* (January 9, 1985).



## 5. COMPUTER NETWORK SECURITY

---

**Dr. Dennis K. Branstad**  
National Bureau of Standards

### Abstract

*Security is one of the critical issues in research computer networks that was identified by Congress for investigation and reporting under Public Law 99-383. This legislation requested a study of the critical problems and current and future options in research computer networks that Congress could use in proposing legislation. This section reports the findings of a panel of computer network security experts convened under the auspices of the Federal Coordinating Council on Science, Engineering and Technology to report on this issue.*

*Before the security objectives could be identified for research computer networks, a set of general objectives for the network were identified as being needed. Since they were not available to the panel at the meeting, a set was assumed. First, a national resource of high-performance computing capability was to be available to universities in all states for performing approved research projects. Second, the capability of the network was to increase as improved technology becomes available. Third, the research network would be used to advance higher performance communication technology.*

### 5.1. NETWORK SECURITY OBJECTIVES

In addition to the assumed general objectives, the following security objectives were identified for computer research networks. First, the integrity of scientific research data must be protected to assure that it has not been modified in an unauthorized manner. Second, access to the computational capability, the communications capability, and the data in the network (including end systems) must be controlled and provided only to authorized users. Third, access to authorized users must not be denied or delayed by unauthorized activities in the network. Fourth, the confidentiality of unpublished scientific data and unpublished results must be assured. Finally, the security mechanisms that are implemented in the network should minimize performance reduction. The panel felt that any failure of achieving these objectives could result in loss of a valuable national asset, loss of scientific leadership of the nation, or loss of credibility to the scientific community as well as embarrassment to all concerned.

### 5.2. GENERAL OBSERVATIONS

The security of a research network of computers, especially one of supercomputers, is one example of a bigger national issue, i.e., the security of all information and information processing capability of the nation including economic, personal, financial, societal, and national security information. While the focus of the study was on research computer networks, the results of the study could be generalized to many information processing networks. Security must be viewed from an overall information system perspective and not just from a technical perspective related to either communications or computers.

Uncontrolled connectivity among information processing systems increases vulnerabilities of the systems unless adequate administrative, physical, and technical security provisions are utilized. The growth and interconnection of computer networks now make the task of analyzing possible paths from one computer in a network to another often very difficult. Additional connections to the network are designed to be unnoticeable to users of the network and hence new vulnerabilities could also be unnoticeable. A recent set of computer breakins through an international computer network, including numerous subnetworks, has demonstrated many vulnerabilities to computer networks and end systems. Such a network includes many policies, managers, control structures, and equipment providers. Increased access is desirable but comes at a price of requiring increased security or increasing vulnerability.

Security must be an objective of all participants in a computer network, including the policy makers, the designers, the implementors, the managers, the users, and the maintainers. Security will only be as good as the security policy under which the network was designed and is operated. While incremental improvements in security are possible, a security architecture is required that is complete and pervasive while uncumbersome. For the time being, significant investments in research and development are needed in order to achieve a range of security necessary to achieve a range of security goals.

### 5.3. THREATS, VULNERABILITIES, AND RISKS

A threat is an expression of intent to damage or cause an undesirable action to something--a computer network (including end computer systems) in this case. A vulnerability is a condition that makes something (e.g., a computer network) open to attack or accidental events. Risk is the potential loss if a vulnerability is exploited.

Threats to computer research networks include intentional destruction, modification or disclosure of data, and unauthorized use of the computational or communications capability of the network, either by unauthorized people for any purpose or by authorized people for unauthorized purposes. Use of a federally funded research computer network for personal gain through unauthorized personal services is an example of the latter. Use of a supercomputer that cannot be exported to Eastern Block countries by citizens of Eastern Block countries via international communications networks is an example of the former.

Vulnerabilities of research computer networks include loss of confidentiality of data because of passive monitoring of communications, especially satellite links; loss of integrity because of active modification of data in network gateways or storage facilities; or denial of service to authorized users by unauthorized use of communications or computational capacity. Vulnerabilities exist if electrical power and communications capabilities have not been assured with adequate back-up facilities. Data is vulnerable to loss if adequate back-up storage is not provided. Other vulnerabilities include potential losses because of inadequate user authentication, access control, fire protection, system recovery capabilities, etc.

Risks (i.e., exploitation of vulnerabilities) of research networks may range from loss of revenue due to unauthorized use of the network to loss of national security by the unauthorized performance of weapons research or cryptanalysis by unauthorized individuals representing countries not having access to such resources (e.g., supercomputers). Loss of anonymity is possible if research reviews are provided via electronic mail on a research network without adequate confidentiality protection. Fraud, waste, and abuse are general risks in all networks and are prevalent in networks in which a significant part of the costs are either not accounted for or are paid for by overhead.

## 5.4. RESEARCH COMPUTER NETWORK SECURITY PROGRAM

A computer security program is necessary for any information processing system and must include a security policy, security objectives, security plans, security mechanisms, and security reviews by auditors. An overall security policy is required, which includes a statement of the goals of the security program. A national security policy may be required for all federally funded research computer networks, especially one that contains supercomputers manufactured in the United States and subject to export controls for national security reasons. Conservation of a national resource requires a national policy. Computer security policies are also required for each host computer (i.e., end system) in the network that incorporates the national network policy and additional policies specific to that system. Examples of clauses that could be in a security policy are contained in Section 6.8. A security program of a large computer network should be able to support and enforce several security policies, including provisions for dynamically changing the policies depending on changing environments or changing emergency requirements.

From the list of possible security services that could be provided in a research computer network, the following security services were considered for research and implementation in a research computer network. First, authentication of the users of the networks on an individual basis should be provided. Second, integrity of the information in the end systems and in the communication network should be assured. Third, access to the end systems should be controlled on an individual basis. Fourth, confidentiality of private data in the end systems and the network should be assured at a reasonable level.

Computer security includes three major categories of protection: administrative, physical, and technical. Administrative procedures include providing back-up of data in alternative storage facilities, alternative electrical power, and communications facilities; and training of users in security policies and procedures; auditing of systems and investigating security relevant events; training of operators in observing security practices and other means of instilling human awareness of good security practices. Physical protection includes mechanical means of preventing unauthorized physical access to computers, communications media, and ancillary equipment. Technical protection includes means of separating users and users data by automated systems in the computer itself through cryptographic systems, "trusted" operating systems, personal authentication systems, and access control systems.

A computer security program can be established for national research computer networks in the near term if a policy is established requiring such a program and resources are provided for implementing such a program. The level of security that is possible in the near term depends on many aspects of the network, especially acceptance by the user community. Administrative and physical protection means can be implemented in the near term and may have already been implemented in many networks. The rest of this report identifies areas of technical protection that, while available in some systems and provide some level of protection assurance, still require significant research in order to provide high levels of assurance that will be required in many networks in the future.

## 5.5. RESEARCH TOPICS

The following research topics were identified by the panel as areas that need more investigation in order to achieve the security goals anticipated for a high-performance research computer network. The topics are provided in an order that corresponds to the four security services identified earlier in the report and not in any order of priority.

### 5.5.1. Personal Authentication

The security panel identified automated personal authentication as a fundamental requirement for supporting a policy of personal accountability in a research computer network. Personal Identification Numbers (PINs) or passwords are typically used for personal authentication in current computer systems. While PINs and passwords are inexpensive methods of personal authentication and can provide reasonable levels of assurance if properly managed, improved methods of personal authentication are desirable and require additional research.

Improved methods of personal authentication involve recognizing tokens (e.g., badges, credit cards, pictures, "smart" cards, etc.) that have been issued to a person and that are used to authenticate the identity of the person to a computer. Alternatively, the computer can be used to measure or recognize a human characteristic that is unique or relatively unique to an individual (e.g., fingerprints, voice, retina patterns). Research is needed in both of these areas in order to authenticate a person with a high level of assurance but without undue expense and inconvenience.

### 5.5.2. Access Control

Access controls to computer networks can be implemented in the end systems (e.g., supercomputers), in the network systems (e.g., terminal controllers), or in special network system (e.g., network access controllers), or in some combination of these. Research is required in how to support any of these approaches with a high level of assurance but without undue expense and inconvenience. Many tradeoffs exist in the approaches and a general network security architecture should be investigated in order to provide a foundation for the research and for development of future networks.

### 5.5.3. Integrity Assurance

Research is necessary in various alternatives for assuring integrity of information during communication, storage, and processing. Technology exists for detecting changes in data that are accidentally or intentionally induced in data during communication or while it is in static storage. For example, cryptographic-based electronic seals have been developed that provide high assurance that a message (e.g., Electronic Funds Transfer message) has not been changed since it was sealed. However, the management of the parameters of the required system, such as cryptographic keys, trusted third parties, trusted implementations, etc., require additional research.

Research is especially required in assuring the integrity during input, processing, and output. Data must be created, changed, and deleted during normal data processing by authorized users. Supercomputers must do these transformations very fast in order to achieve performance goals. Assuring the integrity of these transformations while maintaining high performance requires a significant amount of research.

### 5.5.4. Confidentiality Assurance

Assuring the confidentiality of data during communication, storage, and processing also requires research, especially in a highly vulnerable research computer network. Trusted operating systems have been a subject of research for numerous years and the results of the research are just beginning to pay off. Trusted systems are now being required in some computers and in high-speed communications switches that are necessary to provide high-speed access to the computers. For the most part, the level of confidentiality assurance required is not high in an unclassified research computer network having

numerous students and principal investigators from many states and many countries. However, research is required on how to isolate private data from public data and the private data of one user from private data of other users in a research computer network, especially when access authorization decisions must be made very fast and unobtrusively in order to maintain high-speed processing without inconveniencing legitimate users. Continued research is needed in multilevel security systems for special applications.

### 5.5.5. Security Event Monitoring

Research is required in order to identify what security events require monitoring in a research computer network and how to do the monitoring without invading the rights of researchers and without becoming obtrusive. Security event records can be accumulated much too fast for any human review. Therefore either the records must be reduced, which causes certain security events to be missed, or automated means (e.g., artificial intelligence) must be developed to analyze the events and predict or detect anomalies. Processing of this type requires significant research.

## 5.6. VERIFICATION

The final step in security design is formal verification when high assurance is required in the security services. Formal verification of security means proving mathematically that the security mechanisms of a system are correct and cannot be defeated. This process is very time consuming when performed by humans. Again research is required in automated verification procedures and tools that will minimize the time and effort required to perform security design verification. Each implementation of the design may also have to be verified in order to provide very high levels of security assurance. This is a very long term research goal.

## 5.7. SUMMARY AND CONCLUSIONS

The security panel developed an example set of security objectives of research computer networks and identified that a security policy for the network was the first requirement for a network security program. The threats, vulnerabilities, and risks were then outlined and a hypothetical security policy for a research network was developed. Several near-term security goals that could be achieved using either current technology or the results of ongoing research were identified and some long-term security research activities were proposed. Finally, a general approach to providing security in research computer networks was developed.

The primary conclusions reached are that security has to be a commitment of policy makers, network service providers, research supporters, principal investigators, and users in order to achieve the security goals. Near-term security improvements can be achieved in existing networks if this commitment is made. Research is being performed now in trusted gateway components, trusted computer systems, network security architectures, and trusted terminals, which should provide solutions to some current network security problems. Additional research is required in order to complete the array of necessary security mechanisms that can be used in a variety of networks with a range of security policies, some of which are dynamically changing. Finally, standards are required of several types (e.g., minimum operational standards, interoperability standards, common user interface standards) in order to assure a uniform level of security throughout a widely distributed research computer network.

## 5.8. SECURITY POLICY EXAMPLES: RESEARCH NETWORKS

The following clauses were suggested as examples of those in a security policy without regard to priority, enforcement, or source. These are used as examples and should not be construed as forming a working policy.

- All uses of this research computer network shall be in accordance with the laws of the United States and the laws of the states in which the end systems are located.
- Authorized users of this research computer network are restricted to citizens of those countries that could purchase and use any of the end systems that are attached to or reachable from the network (unless specifically authorized in writing by the designated manager of the network and the designated manager of the end system that is being authorized for use).
- Authorized users of this research computer network shall only perform those activities that are specified by the network manager, end system manager, and principal investigator of the research project if any federal government funds are used to support any activity of that user.
- No user shall perform any act that would cause destruction or modification of any data, including programs, provided by the network or owned by any other user of the network.
- No user shall perform any act that would cause disclosure of information to unauthorized people and that is private to the network or private to any other user of the network.
- Users shall be personally accountable for all activities that they initiate and perform on the network or in any end system of the network.
- Managers shall provide reasonable assurance that all data in the network and end systems of the network are protected against accidental or intentional unauthorized destruction, modification, and disclosure.
- Managers shall provide reasonable assurance that failures and errors in the network and end systems shall cause minimum loss of data and processing capability.
- Users shall assume that their activities in the network and end systems are subject to monitoring and recording for system management and security evaluation purposes.
- Acceptance or assignment of liability for loss of data or processing capability is not assumed by the operators of the network.



## 5.9. WORKING GROUP PARTICIPANTS

Dennis K. Branstad, Chairman  
National Bureau of Standards

Roger Callahan  
National Security Agency

Pete Dillingham  
Cray Computer Research, Inc.

David Golber  
UNISYS Inc.

David Gomberg  
MITRE Corp.

Stephen Kent  
Bolt Baranck and Newman Inc.

Noel Matchett  
Information Security, Inc.

Dan Nessel  
Lawrence Livermore National Laboratories

Douglas Price  
SPARTA, Inc.

George Sullivan  
Defense Communications Agency

Joseph Tardo  
Digital Equipment Corporation

Stephen Walker  
Trusted Information Systems Inc.



## 6. THE GOVERNMENT ROLE IN NETWORKING

---

Jesse H. Poore, Chairman

### Abstract

*Senate Bill S10213 called for a study of "critical problems" and "current and future" options regarding communications networks in the infrastructure of scientific research in the United States. Wording of the bill conveys simultaneous awareness that effective network services for scientists might be approaching limits while, at the same time, the potential for enhanced scientific progress through exploitation of the exciting new technologies of fiber optics, graphics, and supercomputers is enormous. This is indeed the situation, as concluded in an intensive study carried out by some 60 experts in the field. The proper role for government is to structure policies and programs to enable the scientific community to realize the full potential of current technology and to set an agenda for the future.*

### 6.1. INTRODUCTION

The invention of printing in the fifteenth century first enabled systematic scientific communication and exchange, but with time delays of a year or more. Today, computer networking is poised to provide the next major advance in communication and collaboration. Networking supports instantaneous communication, remote collaboration, and remote sharing of expensive research facilities. This encourages the building of research teams on a national and international scale that collaborate on a daily basis rather than just sharing research results. Ideally, scientists geographically dispersed could share a common workbench, whether conducting experiments, writing papers, or using a shared resource such as a supercomputer, accelerator, or large telescope. Ideally, scarce shared resources would be networked into a grid of availability. Whether scientist-to-scientist, laboratory-to-laboratory, or campus-to-campus, a properly devised science network should support connectivity and communication at very high speeds.

The purpose of a national research network is to improve the productivity of the basic research community on solving problems of national interest. The new mode of rapid communication that the network provides will enable more efficient use of expensive national facilities and the building of nationwide research teams with ranges of expertise optimized to solve today's most difficult problems.

### 6.2. SUMMARY

A strong federal role in networking is essential to ensure that a single full function scientific research network emerges from today's Balkanized arrangements. The government's role is prescribed in the following three conclusions.

Conclusion 1: That the National Science Foundation (NSF) be designated the lead agency role in securing interagency cooperation in the full deployment of the nation's scientific and research networks, and that \$120M (FY 1988\$) over the next 5 years be appropriated for the full deployment of current network technology by 1992.

Conclusion 2: That a vigorous program of research and development be initiated to create the basic knowledge and technologies for the next generation network for use by scholars in the United States, and that \$340M (FY 1988\$) over the next 10 years be appropriated for the research and commercialization of research results.

Conclusion 3: That an advanced network for science and industry be designed and deployed by 1997 in a phased program using the new technologies and research results of the above program.

### 6.3. FULL DEPLOYMENT OF CURRENT TECHNOLOGY NETWORKS

Networking in the United States has reached its current state of development through individual agency efforts to satisfy a mission-oriented need, industrial efforts to fill what was perceived to be a market opportunity, or collegial efforts to satisfy the needs of a community of scholars. Examples of the first are the ARPANET, the MFE and HEP networks of the Department of Energy, USAN and NSFnet by NSF, and the networks of NASA and the National Science Board (NSB), NOAA and others. Examples of the second include DECnet, SNA, Ethernet, DOMAIN, Token Ring, and others, both nurtured and abandoned. Examples of the third kind include Edunet, BITNET, and CSNET. These are our success stories, our national exemplars of insightful thinking, daring engineering, entrepreneurship, and community action at its best.

There are problems. An individual scientist may find it necessary to interact with other scientists or machines on more than one of these networks. The networks are usually not compatible because they were developed according to design goals that did not include consideration of uses and technologies unrelated to the job at hand. Each of these networks tends to be especially good at one thing and not so good at another. Some of the networks are overloaded with traffic. Security is not uniformly good from network to network or from host to host. One comes to the question, why not create a new network that will satisfy all needs and resolve the seemingly trivial problems?

The problems are not trivial. Each mission-oriented agency has an accumulated dependency on the technology it uses and, in general, cannot easily switch to an alternative network. It is important that the supercomputer centers be allowed to succeed as measured by their ability to provide computer service to their strategic users, rather than as measured by the means of access. Community networks have a large investment in their support structures. Industry has an interest in both proprietary products and compatibility. So even if all parties are eager to cooperate it is not possible significantly and quickly to change the variety of networks that must continue to exist and grow.

Moreover, current technologies are inadequate for the task. As will be illustrated below, the limits of the current technologies are in sight and we are separated from a new generation of network concepts by several years of basic research.

Thus, the solution in the short term is to press existing technologies for networking as far as possible by dint of clever engineering and careful cooperation, and hope that the new concepts and new technologies will be in hand soon enough.

The first step is to designate a lead agency charged with the responsibility for progress in scientific networks. Upon review of the organic act creating NSF,<sup>1</sup> it is clear that needs just such as this were anticipated by the post-war architects of our peace-time science infrastructure. However, if there are good reasons not to assign this task to an existing agency, one could contemplate a government corporation after the example of the Communications Satellite Corporation (COMSAT).

Interagency cooperation is key to the continued success of all networks and to the rationalization of the networks to the greatest extent possible. A lead agency, with resources to improve connectivity, bandwidth, and support services for all networks is the most direct method within our structure of government. (Consider the Antarctic program as an example of a program in which the NSF has successfully functioned as the lead agency.) The FCCSET committee and its subcommittee on networking would appear to provide a superstructure adequate to facilitate interagency cooperation. A derivative of the current study group might form the kernel of a coordination group at the network planning and implementation level.<sup>2</sup>

NSF must strengthen its own networking program and the resolve to sustain a strong program in order for NSF to have a sufficiently large stake in networking to warrant lead agency status. Statements made in 1985 on behalf of the NSB<sup>3</sup> with respect to this role are clearly adequate if fully implemented. Expectations are very great in the university community for NSF sponsorship of connectivity as well as access to supercomputer centers and other national research facilities. Large investments in networking on university campuses are already being made in response to NSF network initiatives.

It is in the national interest fully to deploy our current technology networks both because we need greater networking capability and because there is not an alternative network to deploy. Full deployment will surely mean different things to different people and its ultimate meaning will come from the lead agency and its working groups. However, a few suggestions are offered which will find easy consensus.

- Increased number of connections. There should be a monotonic increase in the number of scholars given access to the networks used in research.
- Internet connections must be increased. The need for communication among networks must be monitored and adequate gateways must be provided in response to the traffic.
- Coordinated engineering of bandwidth. The several agencies currently purchase bandwidth independently of each other. Coordination of engineering would lead to better distributed bandwidth and more bandwidth for the sums currently expended. A stable funding program is necessary to increase the bandwidth where needed to meet the limits of demand and performance. Formal interagency agreements might be in order to address the sharing of communications lines and the enhancement of capabilities for the common good.
- Maintain reliable service. These networks are now vital to the scientific progress of the country and are clearly beyond the developmental stage insofar as the dependencies of individual scientists are concerned. Consequently, it is necessary to maintain stable and reliable facilities. Internet actions tend to tie together the fates of our networks as well as their media. It is not uncommon for a researcher to pass through three networks and many more gateways, routers, and bridges to achieve the objective. A formal program for assuring internet integrity is required.
- Coordination of network services. The networks vary greatly in the extent of user support services they offer. The lead agency should consider comprehensive directory services and user assistance programs to assure that scientists who cannot get the help they need locally can be directed to a source of assistance.
- Rationalization of current networks. Although it is inappropriate to call for standards in grappling with current networks, certain aspects of certain networks have proved better than competing ideas. Where practical, changes should be made to reduce the incompatibility, to adopt the demonstrably better practices, and to make choices among competing, similar options.
- Incremental improvements in security are possible. However, the tradeoff between security and convenience will be harsh until research in network security produces better tools and techniques. Nevertheless, confidentiality and security of data are essential to most scientists. Resource

protection and restriction to qualified users can be achieved to a significant degree within current technology. Selected transmission paths can be secured against eavesdropping and made available to certain users. A concerted cooperative effort by host system administrators, remote users, and various network administrators (as described by the working group on security) will result in reasonable security at a reasonable cost.

- A structure of accountability for use of the shared network facilities and services should be established. One suggestion for such a structure is to report usage of the shared facilities and services to major aggregate organizations such as regional networks, supercomputer centers, and disciplinary programs, which, in turn, would report in further detail to the campus, laboratory, or scientist as appropriate from case to case. This concept meets the known accountability requirements with minimal effort and overhead.

"Current technology networks" as used herein refers to a collection of topics. In essence, the point is to press the existing concepts as far as possible, realizing that whereas in some aspects the limitations are clear, in other aspects one expects incremental gain to prolong the life of the concept or technology.

Existing agency networks for research are undergoing individual improvement. New networks are actively being planned, budgeted, and built within existing technology but without due regard for improving the general state of affairs. The lead agency should become involved with all such efforts in an attempt to rationalize the activity and let it benefit from existing programs. No unclassified networks should be making changes or coming into being without benefit of FCCSET oversight and lead agency assistance. Otherwise, network diversity will increase and coordination will become ever more difficult.

Private sector companies are offering an ever increasing array of communication services via satellite, recently installed optical fibers, microwave, and reorganized local service. Full advantage should be taken of these offerings as they change from time to time.

The limit of the current technology is very likely not bandwidth or connectivity. Researchers in the field suggest that the limitations will come first in gateways, routers, and switches and then later in the protocols and architectures of the networks. Within full deployment one expects local area networks to reach 300 Mbit/s, campus to wide area network trunks to reach 45 Mbit/s, and connectivity to extend to most campuses of national laboratories, industrial laboratories, and universities. It is clear that limits will be reached in all dimensions.

Full deployment of the capabilities of the existing technology can be achieved by 1991 or 1992. Assuming continuation of the development and planned expansion of existing agency, regional, and discipline-oriented networks, an integrated national research network would be scaled to provide 1.4-Mbit/s transmission rates to about 200 additional institutions. The investment in this network would follow the profile shown in Figure 1 and would reach approximately \$40 to 45M (FY 1988\$) when fully operational. Some existing networks will at first overlap and then transition to the shared facilities with the national research network. Some of the agency networks would continue to operate to provide specialized services in support of certain missions to the 200-300 existing sites.

Sums associated with this deployment cover:

- engineering needs in the early years,
- lease charges of four 45-Mbit/s lines across the country and 1.5-Mbit/s satellite links to remote regions,
- stub lines to the industrial, university, and federal laboratories,
- gateways to the local networks on these sites, and
- management, operation, and maintenance of the national research network.




Figure 1. Costs of an interim network.

#### 6.4. RESEARCH TO DEVELOP THE NEXT GENERATION NETWORK

Future networking needs are expansive in both quantitative and qualitative terms. The goal is to develop the technology that will support the design of a dynamically adaptable, task-configurable network providing to the end user wide area service uniformly comparable to local area service .

One may postulate a community of 1,000 users at each of 1,000 sites connected to the next generation network (as described by the working group on internet concepts). Peak bandwidth requirement on such a network will press the limit of one trillion bits per second, as approximated below.

Increasingly, the results of sophisticated computations are best presented visually. Present examples include fields as diverse as medical imaging and structural analysis. Not only are the results far more enlightening, but the opportunity to interact with computations as they progress offers far more efficient use of both human and computational resources. Unpromising lines of inquiry can be abandoned early while interesting results are pursued in depth. Presenting computer generated images requires 30 frames per second for presentation satisfactory to the human persistence of vision; each image represents about 10 Mbit of information. A reasonable design goal for the national network would be to support one thousand such sessions simultaneously throughout the network, leading to a peak aggregate data rate of 300 Bbit/s.

Architectures of the very highest performance supercomputers tend to be unique as designers try various approaches to push the limits of performance from current technology. Frequently, one computer or architecture will be superior for part of a problem, e.g., numerical computation, and a different machine will be superior for some other aspect, e.g., visual image processing. Exchanging the information

necessary to perform a single computation on multiple, diverse computers at different locations would ideally utilize the full data exchanging capability of each. Currently, such data rates are about 500 to 1,000 Mbit/s. Allowing for some improvement in input-output technology, a transfer rate of 3 Gbit/s between supercomputers is postulated. One hundred such transfers taking place simultaneously throughout the network leads to an aggregate peak rate of 300 Gbit/s to support supercomputer to supercomputer linkages.

Accumulated traffic of more modest description, network management, security systems and other forms of overhead will account for the remaining 40%.

A research program adequate to the goals outlined above will have many dimensions. New network architectures must be explored. Packet switching versus circuit switching must be revisited in light of VLSI technology and other advances. Current network technology is generally 20 years old. New protocols must be developed that better match tasks with pathways. Advances will be needed on all fronts (architectures, components, software, management systems, security, and standards) to meet the challenge.

Research on network and computer security should lead to a technology that delivers its benefits at very low cost in dollars, overhead, and convenience.<sup>4</sup> Areas in which research is expected to be especially fruitful are personal authentication, real-time security event monitors, encryption, and verification tools. Comprehensive network security architecture should allow individual users to invoke an appropriate level of security in terms of encryption, secure paths, and permissions.

Agencies of the federal government have many proven methods of cooperating in research. Certain agencies are better able to undertake aspects of the total program. A measure of overlap and duplication is healthy in large research programs to allow competing ideas to be pursued. Established contract and grant mechanisms could be used to implement this research program; alternatively, a lead agency might be designated and funded for all or part of the program.

Whatever the funding and management mechanism, industry participation is essential for several reasons. Much of the nation's research capability in networking resides in industrial laboratories. Furthermore, the research program envisioned here will best contribute to improve U.S. competitiveness if commercialization is an ever present consideration.

A final concern is the continuing preparation of an adequate number of future scientists and engineers in this field. If the above program of research is initiated in 1990, deployment of the next generation network could be under way by 1997. The network research program must be organic, continually evolving as new techniques and technologies enhance capability. The next generation network must grow out of the research program in planned and coordinated cycles of research, design, prototyping, and testing.

The estimated research investment to bring the capabilities described above to 1,000 users at 1,000 sites is \$340M (FY 1988\$) over the next 10 years, after which time the deployment and operation of the network dominates the continuing research, as described in the next section. The cost profile to allow deployment of the next generation network to begin in 1997 is shown in Figure 2. This estimate is consistent with research and development programs of this scale in large industrial and public works projects--the correct analogy for a program of this type.






Figure 2. Research costs for the next generation national research network.

## 6.5. DEPLOYMENT OF THE NEXT GENERATION NETWORK

An important feature of the world-wide telephone system is that instantaneous voice communication takes place at speeds the human ear expects. Analogously, an important goal for networking is that instantaneous communication of visual data and video takes place at speeds the human eye expects. The next generation network is designed to achieve this goal.

Full deployment of the current technology networks will lead to a large installed base of both network equipment and network users. Researchers in the U.S. will be introduced to a new way of thinking and communicating, both with computers and with each other. A graceful metamorphosis from the old network to the new must be planned. Many lessons have already been learned in network technology and architecture; more will have been learned in the process of fully deploying today's technology in a rational national research network.

The planning process must begin with a thorough understanding of the characteristics and limits of the current networks. Engineering will be needed to understand the limitations in the packet routing systems, which can only be done with a combination of laboratory and full-scale field data tests. Network measurements, improved interfaces, improved gateways, and prototyping in production environments will make the current networks more productive and will guide the conceptualization of the next generation network.

The technology of billions of bits per second data transmission systems will require the extensive development and laboratory testing of new communications hardware, computer interfaces, transmission and routing protocols, and software design. The radically new designs that result will require extensive laboratory and prototype user testing. It is reasonable to expect that these tests will reveal problems to be overcome, so that several cycles of design and prototype testing should be expected.

The outcome of this process should be a final design for a new national research network linking researchers and national support facilities such as supercomputer centers and research instruments. The first phase of deployment would involve settling the network design. Deployment of the trunks would follow, allowing interfaces to individual university campuses and research institutions.

It is important to realize that this deployment is a process and not a project. A national network to support research must be woven into the fabric of the national research infrastructure.

Sums required to deploy the next generation network will depend upon progress in software and hardware development. However, tariff structures created for voice communication are being imposed on data communication. This will overwhelm all other factors. Historically, the capacity to cost ratio based upon voice traffic has halved every 5 years. Communication in the future may be digital rather than analog, and federal and state governments should consider this shift in technology in setting transmission tariffs.

Many assumptions and long extrapolations must be made from the traditional tariff base to estimate the costs of deployment of the next generation network. Of the two major cost components, "line costs" are estimated at \$190M (FY 1988\$) per year, while all other costs (gateways, maintenance, operation, security, and management) are estimated at only \$170M (FY 1988\$) per year. These numbers are clearly unacceptable and one must trust that research will show a way to avoid "line costs," or competition will drive them down by two or three orders of magnitude, or that the basis of calculating "line costs" will change to reflect the new reality of data as a significant market. In the absence of change and progress, the cost of operating the next generation network in the steady state would be \$360M (FY 1988\$) per year.

## 6.6. WORKING GROUP PARTICIPANTS

J. H. Poore, Chairman  
University of Tennessee

Ken Kliwer  
Purdue University

Daniel Drobnis  
San Diego Supercomputer Center

Daniel Masys  
National Institutes of Health

William Johnson  
Digital Equipment Corporation

Abraham Peled  
IBM Watson Research Center

Robert M. Johnson  
Florida State University

Joseph Timko  
AT&T Bell Labs

Robert Kahn  
National Research Initiatives

Kenneth Wilson  
Cornell University

John Killeen  
Livermore NMFEC

## 6.7. REFERENCES

1. National Science Foundation Statutory Authority
2. B.M. Leiner, "Implementation Plan for Interagency Research Internet", *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States* (June 1987).
3. Testimony by M. Good, Committee on Science and Technology, June 1985.
4. R.M. Johnson, "On the Role of the Federal Government in the Management, Operation and Funding of National Research Networking Requirements", *Volume III, A Report to the Congress on Computer Networks to Support Research in the United States* (June 1987).



## GLOSSARY

---

Definitions for the terms in this glossary (other than acronyms) are from the *Computer Dictionary*, Sipple and Sipple, Third Edition, and from *A GLOSSARY for Users of the Integrated Computing Network*, Los Alamos National Laboratory.

### A

ACI	The Advanced Computing Initiative graphics project at Los Alamos.
AI	Acronym for artificial intelligence.
ANSI	American National Standards Institute.
ARPANET	A resource-sharing network sponsored by the Advanced Research Projects Agency.
ASCII	American Standard Code for Information Interchange - an 8 bit code.

### B

BITNET	A resource sharing network sponsored by the participating sites.
bit/s	Bits per second.
broadband	Also called wideband. Pertaining to a channel with a bandwidth greater than voice-grade channels.
BSD	Berkeley Scientific Development
buses	Circuits over which data or power are transmitted.
bytes	Indicates a measurable portion of consecutive binary digits (bits). Common sizes are 6, 8, and 9 bits.
byte/s	Bytes per second.

### C

CCF	Central Computing Facility (Los Alamos).
CCITT	Consultative Committee for International Telephony and Telegraphy.
CC&R	Commitment, concurrency, and recovery.
CFD	Acronym for computational fluid dynamics.
CGM	Computer Graphics Metafile.
CPU	Acronym for central processor unit.
CSLAN	The Computing Services segment of the Lawrence Berkeley Laboratory network.
CSMA	Carrier Sense Multiple Access.

CSNET	A computer science researcher's network that provides a gateway to ARPANET/MILNET.
CSRLAN	The Computer Science Research segment of the Lawrence Berkeley Laboratory network.
CTSS	Cray Timesharing System.

**D**

DACS	Digital Automatic Cross-Connect Systems.
DARPA	Defense Advanced Research Projects Agency.
DCA	Defense Communications Agency.
DDS	Digital Data Service.
DEC	Digital Equipment Corporation.
DECnet	A resource-sharing network of DEC.
DES	Data Encryption Standards
dialup	The service whereby a dial telephone can be used to initiate and effect station-to-station telephone calls.
DOE	Department of Energy.
DoD	Department of Defense.
DP	Acronym for distributed processor.
DR11W	A communications interface component manufactured by DEC.
DS-0, DS-1, DS-3	Data transmission designations representing 64 kbit/s, 1.5 Mbit/s, and 45 Mbit/s data rates.
DSUWG	Data Systems Users Working Group.

**E**

EBDIC code	An eight-level code similar to the American Standard Code for Information Interchange. Abbreviation for expanded binary coded decimal interchange code.
EGP	Exterior Gateway Protocol.
EIN	European Information Network
ESnet	Energy Science Network of the DOE.
Ethernet	A high-speed local area networking technology based in IEEE Standard 802.3

**F**

FCCSET	Federal Coordinating Council on Science, Engineering and Technology.
FDDI	Fiber Distribution Data Interface.
firmware	Computer programs that are embodied in a physical device that can form part of a machine.
FTP	Acronym for file transfer protocol.

**G**

GaAs	Gallium arsenide.
Gbit	Abbreviation for gigabit, which is one billion bits.
Gbyte	Abbreviation for gigabytes.
GGP	Gateway-to-Gateway Protocol.
GKS	Graphical Kernel System.
GSFC	Goddard Space Flight Center.

**H**

HDLC	High-Level Data Link Control.
HEPnet/LEP3NET	Networks that Support High Energy Physics programs of DOE and NSF.
HYPERchannel	A registered trademark of Network Systems Corporation. Defines a high-speed data transmission bus.

**I**

IBM	International Business Machines.
IEEE	Institute of Electrical and Electronic Engineers.
IFIP	International Federation for Information Processing.
IGP	Interior Gateway Protocol.
InteCom	A digital switch from DEC.
IP	Internet Protocol.
IRAB	Internet Research Activities Board.
IRI	Interagency Research Internet.
IRIO	Interagency Research Internet Organization.

IRIS	Integrated Raster Imaging System workstation from Silicon Graphics Inc.
ISDN	Integrated Services Digital Network.
ISO	International Standards Organization.
ISP	Gould Image Support Processor.

**J**

JPL	Jet Propulsion Laboratory.
-----	----------------------------

**K**

kbit	Abbreviation for kilobit. A kilobit is one thousand binary bits.
kbit/s	Abbreviation for kilobits per second.
kbyte	Abbreviation for kilobytes.
kHz	Abbreviation for kilohertz.

**L**

LAN	Local area networks.
LaRC	NASA's Langley Research Center.
LeRC	NASA's Lewis Research Center.
LU 6.2	A general-purpose interprogram synchronous protocol for distributed applications developed by IBM.

**M**

mainframe	The fundamental portion of a computer, i.e., the portion that contains the CPU and control elements of a computer system.
MAN	Metropolitan Area Network
MAP	Manufacturing Automation Protocol
Mbit	Abbreviation for million bits. An Mbit is one million binary bits.
Mbit/s	Abbreviation for million bits per second.
Mbytes	Abbreviation for million bytes per second.
MFENET	A communication network that supports the National Magnetic Fusion Energy Program.



MFLOPS	Million floating-point operations per second.
MHz	Abbreviation for megahertz.
MILNET	A resource-sharing network formerly a portion of ARPANET. The DoD Military Network.
MINET	A resource-sharing European network sponsored by DARPA.
MIPS	Million instructions per second.
MOA	Memo of authorization.
MSC	Mission Support Communications.
ms	Abbreviation for millisecond.
MSFC	Marshall Space Flight Center.
MVS	Memory Virtual System, a trademark of IBM.

## N

NAS	Numerical Aerodynamic Simulation (a NASA program).
NASA	National Aeronautics and Space Administration.
NASnet	A resource-sharing network sponsored by NASA.
NCS	National Communications System
NEC	Nippon Electron Corporation.
NESDIS	National Environmental Satellite and Data Information Services.
NeWS	Network Extensible Window System, trademark of Sun Microsystems, Inc.
NFS	Network File System.
NJE	Refers to a compatible set of IBM software to provide file transfer and electronic mail.
NMFECC	National Magnetic Fusion Energy Computing Center.
NOC	Acronym for network operation center.
NSF	National Science Foundation.
NSFnet	A resource-sharing network sponsored by NSF.
NSI	NASA Science Internet.
NTSC	National Television Standards Code.

**O**

ODA	Office of Document Architecture.
ODC	Acronym for other direct costs.
OER	Office of Energy Research.
online	Descriptive of a system and peripheral equipment or devices in a system in which the operation of such equipment is under control of the CPU.
Op code	A command, usually given in machine language.
OSI	Open Systems Interconnect.
OSTP	Office of Science and Technology Policy.

**P**

PBX	Private branch exchange.
PC	Acronym for personal computers.
PDP11	A minicomputer manufactured by DEC.
PIN	Acronym for personal identification numbers.
PLOT3D	Software used in plotting three-dimensional graphics.
ProNET	A high-speed network implemented at the University of Illinois.
Proteon	Proteon, Inc.
PSC	Program Support Communications.
PSCN	NASA's Program Support Communication Network.

**R**

RGB	Acronym for color television transmission of red, green, blue.
RIACS	Research Institute for Advanced Computer Science.
RJE	Remote Job Entry.
RPC	Remote Procedure Call.

**S**

SDLC	Acronym for synchronous data link control.
SDSC	San Diego Supercomputer Center.

SMTP	A mail protocol.
SNA	System Network Architecture.
SPAN	Space Physics Analysis Network.
SURAnet	A network connecting a consortium of universities throughout the southeastern U.S.

## T

T1	See DS-0, DS-1, DS-3.
TAB	Technical advisory board.
TCP	Transmission Control Protocol.
Telnet	Terminal access network.
TEXnet	Texas Academic Network.
Trans	Transaction Network Service

## U

UCAR	University Corporation for Atmospheric Research.
UNICOS	A UNIX operating system for the Cray computers.
UNIX	A timesharing operating system developed by Bell Laboratories.
USENET	A world-wide distribution network for electronic mail.
USG	Ultra-High-Speed Graphics program at Los Alamos.
UUCP	UNIX to UNIX Communication Protocol.

## V

VAX	A family of mini-computers developed by Digital Equipment Corporation.
VCR	Video Cassette Recorder.
VLSI	Acronym for very large scale integration.
VME	Virtual Memory Extension, a new operating system and a trademark of IBM.
VMS	Virtual Memory System, a trademark of DEC.

**W**

**WDM**                      Wavelength Division Multiplexing.

**X**

**XNS**                      Xerox Network System.

## INDEX

- access control, 60
- accessibility/connectivity, 22
- ANSI, American National Standards Institute, 32
- Antarctic program, 67
- ARC, 31
- ARPANET, 7, 12, 27, 31-32, 34, 66
- AT&T, 7
  
- backbone networks, 33
- bandwidth, 22
- bit rate, 36
- BITNET, 7, 24, 31-32, 66
- bitstreams synchronous, 37
- Bolt Beranek and Newman, 28
  
- capital investment costs, 11
- CERN, 29
- circuit networks, 31
- classes of traffic, 35
- closely coupled supercomputers, 35
- coherent optics, 37
- color graphics, 13
- common carriers, 37
- communication costs, 37
- communication software, 38
- Communications Satellite Corporation (COMSAT), 66
- community users, 35
- competitive advantage, 14
- competitive position, 7
- component failure, 13
- compress data, 11
- computer networks, 31
- computer-generated images, 69
- congestion control, 34
- congestion control algorithms, 11
- consortia networks, 33
- cost, 23
- cost analysis, 13
- cost/capacity forecasts, 10
- CRAY-1s, 29
- cryptographic keys, 60
- CSNET, Computer Science Research Network, 24, 28, 66
- customized network service, 10
  
- DARPA, Defense Advanced Research Projects Agency, 31
- data compression, 13
- Data PBX switches, 30
- Data Systems Users Working Group, (DSUWG), 31
- DECnet, 30, 32, 66
- DECnet Coloured Books protocols X.25, 30
- Defense Advanced Research Projects Agency (DARPA), 27
- Defense Communications Agency (DCA), 23, 27
- Defense Data Network (DDN), 31
- Department of Defense (DoD), 34
- Department of Energy (DOE), 34, 66
- deregulation, 37
- Digital Automatic Cross-Connect Systems (DACS), 37
- digital circuits, 10
- digital cross-connect technology, 34
- digital cross-connect transmission, DACS, 37
- digital data service (DDS), 8
- Digital Equipment Corporation (DEC), 29
- distributed file system, 9
- distributed operating systems, 13
- distributed scientific databases, 13
- distributed window systems, 13
- DoD ISO protocols, 34
- DOE networks, 29
- DOMAIN, 66
- DS-0, 37
- DS-1, 28, 37
- DS-3, 37
  
- Eastern Block countries, 58
- eavesdropping, 68
- economic scientific competitiveness, 21
- Edunet, 66
- Electronic Funds Transfer, 60
- electronic mail, 21
- encryption, 11
- Energy Sciences Network (ESnet), 29
- engineering bandwidth, 67
- Ethernet, 66
- Ethernet long-haul bridges, 31
- European HEP DECnet, 30

- evolutionary development, 12
- expanded coverage, 13
- extensibility, 12
- Fast packet switching, 35
- Federal Coordinating Council for Science, Engineering and Technology (FCCSET), 1
- Fermilab, 29
- Fermilab-CERN, 30
- fiber optics, 8, 10-11, 38
- Frequency Switch, 36
- GaAs, Gallium Arsenide, 37
- gateway computers, 11
- gateway technology, 13
- Genetics Bank, 32
- geographic networks, 33
- global management and planning, 12
- Goddard Space Flight Center (GSFC), 31
- graphics, 22
- GSFC, 31
- HEPnet, 7, 29-30
- hierarchical topologies, 33
- High Energy Physics (HEP) programs, 29
- homogeneity/transparency requirement, 38
- host/mail Name Servers, 34
- ILLIAC IV, 31
- inhomogeneity, 36
- Integrated Services Digital Network (ISDN), 10
- integrated switching, 35
- integrated transmission, 35
- interactive monochrome graphics, 9
- interagency coordinating committee, 7
- interagency sharing, 8
- intercity trunks, 37
- international communications networks, 58
- international sites, 9
- internet DOE, 29
- Internet Protocol (IP), 32
- IP-based protocols, 30
- ISO protocols, 30
- Japan, 29
- Jet Propulsion Laboratory (JPL), 31
- Lawrence Berkeley Laboratory (LBL), 30
- LBL-Japan, 30
- LEP3NET, 29
- LEP3NET transatlantic link, 30
- local area networks, 8
- Marshall Space Flight Center (MSFC), 31
- Massachusetts Institute of Technology (MIT), 30
- media costs, 11
- MFENET, 7, 21, 29
- microwave, 11, 68
- MILNET, 7, 31-32
- mission support communications (MSC), 31
- Mission-oriented networks, 33
- mobile radio, 27
- NASA Networks, 31
- NASF, 31
- NASnet, 7, 21, 31
- National Aeronautics and Space Administration
  - NASA, 23, 34
- national backbone, 28
- National Bureau of Standards (NBS), 34
- national communications infrastructure, 32
- national leadership, 13
- National Library of Medicine, 32
- National Magnetic Fusion Energy Computing Center (NMFEC), 23, 29
- national priority, 27
- National Science Board (NSB), 66-67
- National Science Foundation (NSF), 34, 65
- national security, 21
- network
  - backbones, 8
  - demand, 8
  - design, 13
  - file systems, 13
  - management, 7, 23
  - management authority, 13
  - transparency, 9
  - utilization, 8
- network leadership position, 7
- networking requirements and future alternatives, 7
- nonswitched circuits, 10
- NSF networks, 28
- NSFnet, 7, 28, 32, 34, 66
- Numerical Aerodynamic Simulation (NAS) program, 31

- Office of Energy Research (OER), 29
- Office of Science and Technology Policy (OSTP), 1
- Open Systems Interconnect (OSI), 29
- optical fibers, 37, 68
- optimal network services, 11
- optimum backbone locations, 34
- overloaded gateways, 34
- overloaded networks, 34
  
- packet networks, 8
- PACKET RADIO, 27
- packet switches, 28, 33
- packet switching, 10, 38, 70
- passwords, 60
- performance, 34
- performance degradation, 8
- Personal Identification Numbers (PINs), 60
- Phonet, 28
- Program Support Communications (PSC), 31-32
- protocols, 28
- Public Law 99-383, 1
- public packet networks, 31
  
- Reliability, 12
- Reliability and Recovery, 22
- reliable service, 67
- remote access, 21
- remote resources, 13
- responsiveness, 23
- risk, 58
- risks of research networks, 58
- robustness, 33
- routing algorithms, 28, 33
  
- San Diego Supercomputer Center (SDSC), 23
- satellite, 11, 27, 38, 68
- SATNET, 27
- scientific productivity, 22
- SDSCnet, 21
- security, 68
  - event monitoring, 61
  - for research computer networks, 57
  - policy examples, 62
  - requirements, 22
  - verification, 61
- security/privacy, 22
- sharing resources, 34
  
- SLAC, 29
- SNA, 66
- Space Physics Analysis Network (SPAN), 7, 30-32
- standards, 22
- supercomputer, 14, 21, 28, 33
  - access, 8
  - control of the access, 22
  - network, 22
  - performance, 21
- Superconducting supercollider, 30
- switching technology, 36
- switchover time, 36
- Switzerland, 29
  
- T1, 37
- T1 trunks, 31
- T3, 37
- telecommunications corporations, 10
- Texas Academic Network, TEXnet, 31
- threat, 58
- threats to computer research networks, 58
- Time Division Multiplexing (TDM), 33
- token ring, 66
- topology, 23
- transmission control, 28
- Transmission Control Protocol (TCP), 32
- Transmission Control Protocol/Internet Protocol (TCP/IP), 28, 34
- transparent operation, 22
- trusted implementations, 60
- trusted operating systems, 61
- trusted third parties, 60
- tuned protocol, 12
- tuning existing protocol, 11
- type-of-service routing, 11
  
- U.S.
  - academic community, 32
  - leadership, 21
  - prestige, 9
- University Corporation for Atmospheric Research (UCAR), 28
- upgrading trunks, 11
- USAN, 66
- user services, 12
  
- Video and teleconferencing, 35

video compression, 11  
Vitalink boxes, 31  
voice/data integration, 37  
vulnerabilities of research computer networks, 58

Wavelength-Division Multiplexing (WDM), 37  
WBNET, 27  
West Germany, 29  
wide area network technology, 7  
wide area networks, 8  
Workshop Participants, 4  
Workshop--Agenda, 2  
workstation, 38  
workstation-to-host, 35  
workstations, 33  
world leadership, 11, 14

X.25, 28-29  
X.25NET, 28



A Report to the Congress on  
Computer Networks to Support Research  
In the United States

\*\*\*\*\*

A Study of Critical Problems and Future Options

Volume III  
A Compendium of Supporting Technical Data

Preliminary Issue

June 1987

Document prepared by the  
Computing and Communication Division,  
Los Alamos National Laboratory, Los Alamos, New Mexico.\*

This report was prepared as an account of work sponsored by agencies of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

\*The Los Alamos National Laboratory is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

# CONTENTS

---

## INTRODUCTION TO THE STUDY

## INTERNET CONCEPTS

1. Implementation Plan for Interagency Research Internet
2. The NASA Science Internet
3. The Role of Regional Networks

## NETWORKING TRENDS

4. Networking: Some Observations on the Future
5. Future Directions in Communications Research
6. Advanced Systems Software for Supercomputers
7. Local Area Network Technology with Emphasis on Gateways and Digital Switches
8. The Role of Graphics Workstations in Supercomputing
9. Networking Trends in Data Communications
10. Industry and Technology Trends
11. AT&T Presentation Graphics
12. Contel Presentation Graphics

## NETWORKING REQUIREMENTS

13. The Effect of Distributed Computing Technology on Wide Area Network Requirements
14. The Impact of Distributed Functions on Network Requirements
15. National Network Requirements: A Los Alamos Perspective on Computer Graphics and Distributed Computing in a Supercomputer Environment
16. Network Requirements for Scientific Research
17. Summary of Agency Network Plans and User Community
18. DOE Networking Requirements
19. Research University Bandwidth Estimates

## GOVERNMENT ROLE

20. The Federal Government's Role in National Research Networking
21. Role of the Federal Government in National Research Networking Requirements

APPENDIX A--Additional References

APPENDIX B--Glossary

## CONTENTS

---

CONTENTS .....	v
Introduction to the Study .....	1
<b>1. Implementation Plan for Interagency Research Internet .....</b>	<b>5</b>
1.1. Introduction .....	5
1.1.1. Motivation .....	6
1.1.2. Summary of Previous Report .....	6
1.1.3. Overview of Implementation Plan .....	7
1.2. Background .....	7
1.2.1. History of Research Networks .....	8
1.2.2. Internet Model .....	9
1.3. Technical Approach .....	11
1.3.1. Communications Infrastructure .....	11
1.3.2. User Services .....	12
1.3.3. Internet Research Coordination .....	12
1.4. Management Approach .....	12
1.4.1. General Management Structure .....	13
1.4.2. Funding .....	15
1.4.2.1. Operations Costs .....	15
1.4.2.2. Research Costs .....	16
1.5. Phased Implementation Plan .....	16
1.5.1. Technical Phasing .....	16
1.5.1.1. Phase I--Functional Interoperability .....	17
1.5.1.2. Phase II--Full IRI Capability .....	17
1.5.2. Administrative Phasing .....	18
1.6. Industry Role .....	18
1.7. Funding Breakdown .....	19
1.8. Summary and Conclusions .....	19
1.9. Acknowledgment .....	20
1.10. References .....	20
<b>2. The NASA Science Internet .....</b>	<b>23</b>
<b>3. The Role of Regional Networks .....</b>	<b>25</b>
<b>4. Networking: Some Observations on the Future .....</b>	<b>29</b>
4.1. Introduction .....	29
4.2. Networks as Computers .....	30
4.3. Moving Processes Instead of Data .....	31
<b>5. Future Directions in Communications Research .....</b>	<b>35</b>
5.1. Technology .....	35
5.1.1. The Future of Computing Networks .....	36
5.1.2. Supercomputers .....	36
5.1.3. Parallel Computing .....	36

5.1.4. Artificial Intelligence .....	37
5.1.5. Teleconferencing .....	38
5.1.6. Image Processing .....	38
5.1.7. Digital Document Distribution .....	38
5.2. Future Network Architectures .....	39
5.2.1. Network Size Estimates .....	39
5.2.2. Network Speed Estimates .....	39
5.2.3. Resource Management .....	40
5.2.4. Addressing Dynamics .....	40
5.2.5. Organization Structure .....	41
5.2.6. Data Routing .....	41
5.2.7. Network Security .....	41
5.3. Projects for Today .....	42
5.3.1. New Switching Facilities .....	42
5.3.2. Large Networks .....	42
5.3.3. Network Management .....	43
5.3.4. Multiorganizational Networks .....	43
5.4. Summary .....	43
<b>6. Advanced System Software for Supercomputers .....</b>	<b>45</b>
6.1. Introduction .....	45
6.2. Advanced Systems Software Requirements .....	45
6.3. Distributed Computing .....	46
6.3.1. User Requirements .....	46
6.3.2. Systems Requirements .....	46
6.4. Multiprocessing .....	47
6.5. System Software .....	48
6.6. Requirements for Advanced Systems Software .....	48
6.7. Distributed Computing Requirements .....	49
6.8. Multiprocessing Requirements .....	50
6.9. Conclusion .....	52
6.9.1. Distributed Computing .....	52
6.9.2. Multiprocessing .....	53
6.10. Interprocess Communication Architectures for Distributed Operating Systems .....	53
6.11. Problems with UNIX as a Distributed Operating System for Supercomputers .....	54
6.12. References .....	55
<b>7. Local Area Network Technology with Emphasis on Gateways and Digital Switches .....</b>	<b>57</b>
7.1. Network Functional Requirements .....	57
7.1.1. Electronic Mail .....	58
7.1.2. File Transfer .....	58
7.1.3. Terminal Access .....	59
7.1.4. Distributed Computation .....	59
7.2. Local Area Network Protocols .....	60
7.2.1. DECnet .....	60

7.2.2. TCP/IP .....	60
7.2.3. IBM SNA and NJE .....	60
7.3. Integration of Local Area Networks On Campus .....	61
7.4. Local Area Network Communication Backbones .....	61
7.4.1. Transmission Media .....	61
7.4.2. Network Topology .....	62
7.5. Access Methods .....	62
7.6. Gateway Technology .....	63
7.6.1. NJE Gateways .....	64
7.6.2. DECnet Gateways .....	64
7.6.3. Internet Gateways .....	64
7.6.4. Protocol Conversion Gateways .....	65
7.7. Network Performance .....	66
7.8. Impact of the Digital Switch .....	66
7.9. Network Addressing .....	67
7.9.1. DECnet Addressing .....	67
7.9.2. Internet Addressing .....	68
7.10. Name Servers .....	68
7.11. Network Security .....	69
7.12. Interactions between Campus and Wide Area Networks .....	70
7.13. Topics for Future Study .....	71
7.14. Conclusions .....	71
7.15. References .....	71
<b>8. The Role of Graphics Workstations in Supercomputing .....</b>	<b>73</b>
8.1. Overview of the NAS System .....	73
8.2. Workstation Requirements for the NAS Program .....	73
8.3. Description of Present Workstation Hardware .....	75
8.4. The Network Environment .....	75
8.5. The Cray-2 as a Graphics Co-processor .....	76
8.6. Concluding Remarks .....	78
8.7. References .....	79
<b>9. Networking Trends in Data Communications .....</b>	<b>81</b>
9.1. Classification of Data Transmission Types .....	81
9.2. Present Networking Approach .....	82
9.3. Emerging Networking Approach .....	83
9.4. Trends in Carrier Services .....	85
9.5. Network Standards .....	85
9.6. Conclusions .....	86
<b>10. Industry and Technology Trends .....</b>	<b>87</b>
<b>11. AT&amp;T Presentation Graphics .....</b>	<b>89</b>
<b>12. Contel Presentation Graphics .....</b>	<b>101</b>
<b>13. Effect of Distributed Computing on Wide Area Network Requirements .....</b>	<b>121</b>
13.1. Background .....	121

13.2. Distributed Computing Products--A Review .....	122
13.3. Why Users Need Local Area Network Technology on Wide Area Networks .....	123
13.4. A Comparison of Traditional and Modern Distributed Computing Traffic .....	125
13.5. Conclusions .....	126
13.6. Distributed Computing Traffic at Lawrence Berkeley Laboratory .....	126
13.7. References .....	130
<b>14. Impact of Distributed Functions on Network Requirements .....</b>	<b>132</b>
14.1. Introduction .....	132
14.2. Why Distributed Processing .....	132
14.3. Specific Distributed Applications .....	133
14.4. Long-Range Distributed Applications .....	134
14.5. Building Blocks and Tools for Distributed Applications .....	135
14.6. Networking Requirements .....	136
14.7. Research Questions .....	138
14.8. Summary .....	139
<b>15. National Network Requirements: A Los Alamos Perspective on Computer Graphics and Distributed Computing in a Supercomputer Environment .....</b>	<b>140</b>
15.1. Los Alamos Computing Environment .....	140
15.2. Productivity .....	142
15.3. Computer Graphics .....	143
15.4. Distributed Processing .....	145
15.5. Requirements of a National Network .....	146
15.6. Summary .....	147
15.7. References .....	148
<b>16. Network Requirements for Scientific Research .....</b>	<b>150</b>
16.1. Introduction .....	150
16.2. Network Functions .....	150
16.2.1. Near-Term Requirements .....	151
16.2.1.1. User/Resource Connectivity .....	151
16.2.1.2. Home Usage .....	151
16.2.1.3. Charging .....	151
16.2.1.4. Applications .....	151
16.2.1.5. Capabilities .....	152
16.2.1.6. Standards .....	152
16.2.2. Long-Term Goals .....	153
16.3. Network Connectivity .....	153
16.3.1. Near-Term Objectives .....	153
16.3.1.1. Discipline Connectivity .....	154
16.3.1.2. Communication Resource Sharing .....	154
16.3.1.3. Expansion of Physical Structure .....	154
16.3.1.4. Network Engineering .....	155
16.3.1.5. Support of Different Types of Service .....	155
16.3.2. Future Goals .....	155

16.3.2.1. Bandwidth .....	155
16.3.2.2. Expressway Routing .....	156
16.4. Network Specifications .....	156
16.4.1. Near-Term Specifications .....	156
16.4.1.1. Interoperability .....	156
16.4.1.2. Access Control .....	156
16.4.1.3. Privacy .....	157
16.4.1.4. Accounting .....	157
16.4.1.5. Type of Service Routing .....	157
16.4.1.6. Administration of Address Space .....	157
16.4.1.7. Remote Procedure Call Libraries .....	157
16.4.1.8. Remote Job Entry (RJE) .....	158
16.4.1.9. Multiple Virtual Connections .....	158
16.4.1.10. Network Operation and Management Tools .....	158
16.4.2. Future Goals .....	158
16.4.2.1. Standards and Interface Abstractions .....	159
16.4.2.2. OSI Transition and Enhancements .....	159
16.4.2.3. Processor Count Expansion .....	159
16.4.2.4. Dynamic Binding of Names to Addresses .....	159
16.5. User Services .....	160
16.5.1. Near-Term Requirements .....	160
16.5.1.1. Directory-Services .....	160
16.5.1.2. Information Sharing Services .....	160
16.5.1.3. Resource Sharing Services .....	161
16.5.1.4. Legalities of Interagency Research Internet .....	161
16.5.1.5. Standards .....	161
16.5.1.6. User-Oriented Documentation .....	162
16.5.2. Future Goals .....	162
16.6. Conclusion .....	162
16.7. Internet Task Force on Scientific Computing .....	163
<b>17. Summary of Agency Network Plans and User Community .....</b>	<b>164</b>
17.1. National Science Foundation Affiliated Networks .....	164
17.1.1. NSFnet Backbone Network and Connections to the Research Community .....	164
17.1.2. NCAR Supercomputer Center .....	164
17.1.3. Remote Job Entry Capability .....	165
17.1.3.1. UNINET .....	165
17.1.3.2. USAN .....	165
17.1.3.3. UNIDATA .....	165
17.1.3.4. CSNET .....	166
17.1.4. International Gateway Usage .....	167
17.2. Department of Energy (DOE) .....	168
17.2.1. Energy Sciences Network .....	168
17.2.2. High-Energy Physics Net (HEPnet) .....	168



17.3. National Aeronautics and Space Administration (NASA) .....	168
17.3.1. NASA Aerodynamics Simulation (NAS) .....	168
17.3.2. Program Support Communication Network (PSCN) .....	168
17.4. National Oceanic and Atmospheric Administration (NOAA) .....	169
17.5. U.S. Navy .....	169
17.6. National Institutes of Health (NIH) .....	169
<b>18. DOE Networking Requirements .....</b>	<b>170</b>
18.1. The DOE Research Environment .....	170
18.2. Existing DOE Network Usage .....	170
18.3. Fundamental Requirements .....	171
18.4. Data Bandwidth Required .....	172
18.5. Current Plans and Projections .....	173
18.6. Partial List of DOE Research Sites .....	174
18.7. Network Maps and Projected Growth Data .....	175
<b>19. Research University Bandwidth Estimates .....</b>	<b>180</b>
<b>20. The Federal Government's Role in National Research Networking .....</b>	<b>184</b>
20.1. Definitions, Functions, and Goals .....	184
20.2. Assuring Success .....	185
20.3. The Federal Role in Network Research .....	185
20.4. Management and Operation of National Research Networks .....	187
20.5. Conclusions .....	188
<b>21. The Role of the Federal Government in National Research Network Requirements .....</b>	<b>190</b>
21.1. Background .....	190
21.2. Present Needs .....	191
21.3. Future Needs and the Role of Federal Government .....	191
<b>Appendix A -- Additional References .....</b>	<b>194</b>
<b>Appendix B -- Glossary .....</b>	<b>196</b>

---

## INTRODUCTION TO THE STUDY

## INTRODUCTION TO THE STUDY

---

The 99th Congress charged the Office of Science and Technology Policy (OSTP) with conducting "a study of critical problems and current and future options regarding communications networks for research computers, including supercomputers, at universities and federal research facilities in the United States" (Public Law 99-383, August 21, 1986). At OSTP's direction, an interagency group under the auspices of the Federal Coordinating Council for Science, Engineering and Technology (FCCSET) was formed to carry out the study.

The Congress asked that the following issues be included in the study:

- the networking needs of the nation's academic and federal research computer programs, including supercomputer programs, over the next 15 years, including requirements in terms of volume of data, reliability of transmission, software compatibility, graphics capabilities, and transmission security;
- the benefits and opportunities that an improved computer network would offer for electronic mail, file transfer, and remote access and communications; and
- the networking options available for linking academic and research computers, including supercomputers, with a particular emphasis on the advantages and disadvantages of fiber optic systems.

Networks are essential components of science, engineering, and technology, making it possible for researchers to share resources, access remote facilities, exchange information and software, and improve their collaborative relationships. The challenge is to make effective use of existing computer networks and to create advanced capabilities for the future. This will not only enhance scientific research in many disciplines, but will also help to advance the state-of-the-art of U.S. networking technology.

In conducting the network study during late 1986 and early 1987, the FCCSET group enlisted the help of many experts from government, industry, and academia. White papers were invited on networking trends, requirements, concepts, applications, and plans. A workshop involving nearly 100 researchers, network users, network suppliers, and policy officials was held in February 1987 to air ideas, gather information, and develop the foundation for the report to the Congress. Industry organizations were invited to provide information on the expected costs and benefits of fiber optic systems for networks.

As a result of the collaborative efforts of many dedicated people, the report to the Congress has been completed. It is published in three volumes:

- Volume I contains the FCCSET recommendations on developing computer networks to support research in the U.S.
- Volume II contains summaries of the February 1987 workshop discussions, which focused on six topics: access requirements and future alternatives, special requirements for supercomputer networks, internetwork concepts, future standards and services requirements, security issues, and the government role in networking.
- Volume III contains the invited white papers.

The workshop summaries in Volume II and the white papers in Volume III are presented as developed by their authors. No attempt has been made to achieve unanimity of opinion; there are many points of view expressed on a variety of network-related subjects.

I gratefully acknowledge the participation and support of the many people who have contributed to this report--the workshop participants; the chairs of the workshop groups; the San Diego Supercomputer Center, which hosted the workshop; the authors of the white papers; the staff of the Los Alamos National Laboratory who prepared the three volumes of the *Report to the Congress on Computer Networks to Support Research in the U.S.*; and the members of the FCCSET group that conducted the study.

My special thanks to John Cavallini for his support to me in organizing this study and preparing the report.

James H. Burrows  
Chair  
Computer Network Study

#### Computer Network Study Group

Ron Bailey, National Aeronautics and Space Administration  
Gordon Bell, National Science Foundation  
James Burrows, National Bureau of Standards  
John Cavallini, Department of Energy  
Michael Corrigan, Department of Defense  
Paul Huray, Office of Science and Technology Policy  
Thomas Kitchens, Department of Energy  
James Oberthaler, Department of Health and Human Services  
Dennis Perry, Advanced Defense Research Projects Agency  
Shirley Radack, National Bureau of Standards  
Rudi Saenger, Naval Research Laboratory  
Daniel Van Belleghem, National Science Foundation  
Stephen Wolff, National Science Foundation

## INTERNET CONCEPTS



# 1. IMPLEMENTATION PLAN FOR INTERAGENCY RESEARCH INTERNET

---

**Barry M. Leiner**

Research Institute for Advanced Computer Science\*  
NASA Ames Research Center

## Abstract

*Networking has become widespread in the scientific community, and even more so in the computer science community. There are networks being supported by a number of the federal agencies interested in scientific research, and many scientists throughout the country have access to one or more of these networks. Furthermore, there are many resources (such as supercomputers) that are accessible via these networks.*

*While many of these networks are interconnected on an informal basis, there is currently no consistent mechanism to allow sharing of the networking resources. Recognizing this problem, the Federal Coordinating Council on Science, Engineering and Technology Committee on Very High Performance Computing formed a Network Working Group in 1985. This group has recommended an administrative and management structure for interconnecting the current and planned agency networks supporting research. The structure is based on the concept of a network of networks using standard networking protocols.*

*This paper elaborates on the earlier recommendation and provides an implementation plan. It addresses three major areas: communications infrastructure, user support, and ongoing research. A management and administrative structure is recommended for each area; and a budgetary estimate provided. A phased approach for implementation is suggested that will quickly provide interconnection and lead to the full performance and functionality as the required technologies are developed and installed. While this report addresses the interconnection of agency networks, and cooperation by certain federal agencies, some discussion is presented of the possible role that industry can play in support and use of such a network.*

## 1.1. INTRODUCTION

Computer networks are critical in providing scientists access to computing resources (such as supercomputers) and permitting computer-supported interaction between researchers. Several agencies, recognizing this need, have established networks to provide the needed communications infrastructure. The need for this infrastructure, though, cuts across the various agencies. To that end, the Federal Coordinating Council on Science, Engineering and Technology (FCCSET) Committee on Very High Performance Computing Network Working Group has recommended the formation of an Interagency Research Internet (IRI).<sup>1</sup>

---

\* Work reported herein was supported by Cooperative Agreement NCC 2-387 from the National Aeronautics and Space Administration to the Universities Space Research Association.

The purpose of this paper is to suggest an implementation plan for such an IRI. It addresses three major areas: communications infrastructure, user support, and ongoing research. A management and administrative structure is recommended for each area, and a budgetary estimate is provided. A phased approach for implementation is suggested that will quickly provide interconnection and lead to full performance and functionality as the required technologies are developed and installed. Finally, some discussion is presented on a possible role for industry in supporting and using such a network.

### 1.1.1. Motivation

The prime responsibility for providing the required infrastructure for successful research lies with the researcher, his/her institution, and the agency supporting that research. Thus, the individual agencies have installed and are continuing to enhance computer networks to allow their researchers to access advanced computing resources such as supercomputers as well as being able to communicate with each other via such facilities as electronic mail.

However, there are a number of reasons why it is advantageous to interconnect the various agency networks in a coherent manner so as to provide a common "virtual" network supporting research.

- a. The need to make effective use of available networks without unnecessary duplication. The agencies each support researchers in many parts of the country, and have installed equally widespread resources. Often, it is more effective for a scientist to be provided networking service through a different agency network than the one funding his research. For example, suppose several scientists at an institution are already being funded by the National Aeronautics and Space Administration (NASA) and are connected to a NASA-supported network. Now a scientist at the same institution but supported by the National Science Foundation (NSF) needs access to an NSF supercomputer. It is much more effective to provide that connectivity through an interconnection of NASA and NSF networks than to establish another connection (to NSFnet) to the same university.
- b. The need to establish communication infrastructure to permit scientists to access resources without regard to which network they are connected but without violating access controls on either the networks or the resources. A scientist may be supported by multiple agencies, and therefore have access to resources provided by several agencies. It is not cost effective to have to provide a separate network connection to the scientist for each of those agency resources.
- c. The need for a communications infrastructure to encourage collaborative scientific research. One of the primary functions of a computer network supporting science is encouraging collaboration between researchers. Scientific disciplines typically cut across many different agencies. Thus, support of this collaboration should be without regard to agency affiliation or support of the scientists involved.
- d. The need for a cooperative research and development program to evolve and enhance the IRI and its components where appropriate. Scientific research is highly demanding of both the computing and networking environment. To assure that these needs continue to be met, it is necessary to continually advance the state of the art in networking, and apply the results to the research networks. No individual agency can afford to support the required research alone, nor is it desirable to have inordinate duplication of research.

### 1.1.2. Summary of Previous Report

The above reasons led to the formation of the FCCSET Committee on Very High Performance Computing and its Network Working Group. This group began in early 1985 to discuss the possibility



of interconnecting into a common networking facility the various agency networks supporting scientific research. These discussions led to the report issued in 1986 recommending such an approach.<sup>1</sup>

The report used the "network of networks" or internet model of interconnection. Using a standard set of protocols, the various networks can be connected to provide a common set of user services across heterogeneous networks and heterogeneous host computers.<sup>2,3,4</sup> This approach is discussed further in Section 1.2.

The report goes on to recommend an administrative and management structure that matches the technical approach. Each agency would continue to manage and administer its individual networks. An interagency body would provide direction to a selected organization who would provide the management and operation of the interconnections of the networks and the common user services provided over the network. This selected organization would also provide for coordination of research activities, needed developments, and reflecting research community requirements into the national and international standards activities.

### 1.1.3. Overview of Implementation Plan

The general structure of the proposed IRI is analogous to a federal approach. Each of the agencies is responsible for operating its own networks and satisfying its users' requirements. The IRI provides the interconnecting infrastructure to permit the users on one network to access resources or users on other networks. The IRI also provides a set of standards and services that the individual agencies, networks, and user communities can exploit in providing capabilities to their individual users. The management structure, likewise, provides a mechanism by which the individual agencies can cooperate without interfering with the agencies' individual authorities or responsibilities.

In this paper, an implementation plan for the IRI is proposed. First, some background is given of the previous efforts to provide networks in support of research, and the genesis of those networks. A description of the suggested approach to attaining an IRI is then given. This description is divided into two sections: technical and management. The technical approach consists of two components. First is the provision of an underlying communications infrastructure; i.e., a means for providing connectivity between the various computers and workstations. Second is provision of the means for users to make effective use of that infrastructure in support of their research. The management section elaborates on the suggestions made in the FCCSET committee report. A structure is suggested that allows the various agencies to cooperate in the operations, maintenance, engineering, and research activities required for the IRI. This structure also provides the necessary mechanisms for the scientific research community to provide input with respect to requirements and approaches.

Finally, a phased implementation plan is presented that would allow the IRI to be put in place rapidly with modest funding. A budgetary estimate is also provided.

## 1.2. BACKGROUND

The combination of packet switched computer networks, internetworking to allow heterogeneous computers to communicate over heterogeneous networks, the widespread use of local area networks, and the availability of workstations and supercomputers has given rise to the opportunity to provide greatly improved computing capabilities to science and engineering. This is the major motivation behind the IRI.

### 1.2.1. History of Research Networks

The Defense Advanced Research Projects Agency (DARPA) developed the concept of packet switching beginning in the mid 1960s. Beginning with the ARPANET (the world's first packet switched network),<sup>5</sup> a number of networks have been developed. These have included packet satellite networks,<sup>6,7</sup> packet radio networks,<sup>8,7</sup> and local area networks.<sup>9</sup>

Although the original motivation for the ARPANET development was computer resource sharing, it was apparent early on that a major use of such networks would be for access to computer resources and interaction between users.<sup>10</sup> Following the ARPANET development, a number of other networks have been developed and used to provide both of these functions.<sup>11</sup> CSNET was initiated to provide communications between computer science researchers.<sup>12,13</sup> CSNET was initiated by the NSF in cooperation with a number of universities, but is now self-sufficient. Its subscribers include universities throughout the world as well as industrial members interested in interacting with computer scientists.

CSNET makes use of a number of networking technologies including the ARPANET, public x.25 networks, and dialup connections over phone lines to support electronic mail and other networking functions. In addition to the basic data transport service, CSNET and ARPANET operate network information centers that provide help to users of the network as well as a number of services including a listing of users with their mail addresses (white pages) and a repository where relevant documents are stored and can be retrieved.

With the installation of supercomputers came the desire to provide network access for researchers. One of the early networks to provide this capability was MFEnet.<sup>11</sup> It was established in the early 1970s to provide Department of Energy (DOE)-supported users access to supercomputers, particularly a CRAY-1 at Lawrence Livermore National Laboratories. Because MFEnet was established prior to widespread adoption of the Transmission Control Protocol/Internet Protocol (TCP/IP) protocol suite (to be discussed below), the MFEnet uses a different set of protocols. However, interfaces have been developed between the MFEnet and other networks, and a migration plan is currently under development.

NASA Ames Research Center has long been in the forefront of using advanced computers to support scientific research. The latest computing facility, the Numerical Aerodynamic Simulator (NAS), uses a CRAY-2 and other machines along with a number of networking technologies to provide support to computational fluid dynamics researchers.<sup>14</sup> This system uses the TCP/IP protocol suite both locally and remotely and provides easy access through advanced workstations.

Recognizing the importance of advanced computers in carrying out scientific research, NSF in 1984 embarked on an ambitious program to provide supercomputer access to researchers. This program involved both the provision of supercomputers themselves (through purchase of computer time initially, and establishment of supercomputer centers) and provision of access to those supercomputers through an extensive networking program, NSFnet.<sup>15</sup> The NSFnet uses a number of existing networks (e.g., ARPANET, BITNET, MFEnet) and exploratory networks interconnected using the TCP/IP protocol suite (discussed below) to permit scientists widespread access to the supercomputer centers and each other. The NSFnet is also taking advantage of the widespread installation of campus and regional networks to achieve this connectivity in a cost effective manner.

The above are only a small number of the current and existing networks being used to support research. Hoskins and Quarterman provide a good synopsis of the networks currently in operation.<sup>11</sup> It is obvious from this that effective interconnection of the networks can provide cost-efficient and reliable services.

Starting in the early 1970s, recognizing that the military had a need to interconnect various networks (such as packet radio for mobile operation with long-line networks like the ARPANET), DARPA initiated the development of the internet technologies.<sup>16</sup> Beginning with the development of the protocols for interconnection and reliable transport (TCP/IP), the program has developed methods for providing electronic mail, remote login, file transfer, and similar functions between differing computers over dissimilar networks.<sup>3,4</sup> Today, using that technology, thousands of computers are able to communicate with each other over a "virtual network" of approximately 200 networks using a common set of protocols. The concepts developed are being used in the reference model and protocols of the Open Systems Interconnection (OSI) model being developed by the International Standards Organization (ISO).<sup>17</sup> This is becoming even more important with the widespread use of local area networks. As institutions install their own networks, and need to establish communications with computers at other sites, it is important to have a common set of protocols and a means for interconnecting the local networks to wide area networks.

### 1.2.2. Internet Model

The DARPA internet system uses a naming and addressing protocol, called the IP, to interconnect networks into a single virtual network. Figure 1 shows the interconnection of a variety of networks into the internet system. The naming and addressing structure allows any computer on any network to address in a uniform manner any computer on any other network. Special processors, called gateways, are installed at the interfaces between two or more networks and provide both routing among the various networks as well as the appropriate translation from internet addresses to the address required for the attached networks. Thus, packets of data can flow between computers on the internet.

Because of the possibility of packet loss or errors, the TCP is used above the IP to provide for reliability and sequencing. TCP together with IP and the various networks and gateways then provides for reliable and ordered delivery of data between computers. A variety of functions can use this connection to provide service to the users. A summary of the functions provided by the current internet system is given in Ref. 4.

To assure interoperability between military users of the system, the Office of the Secretary of Defense mandated the use of the TCP/IP protocol suite wherever there is a need for interoperable packet switched communications. This led to the standardization of the protocols.<sup>18-22</sup>

Thus, the TCP/IP protocol suite and associated mechanisms (e.g., gateways) provide a way to interconnect heterogeneous computers on heterogeneous networks. Routing and addressing functions are taken care of automatically and transparently to the users. The ISO is currently developing a set of standards for interconnection that are very similar in function to the DARPA-developed technologies. Although ISO is making great strides, and the National Bureau of Standards is working with a set of manufacturers to develop and demonstrate these standards, the TCP/IP protocol suite still represents the most available and tested technology for interconnection of computers and networks. It is for that reason that several agencies/programs, including the Department of Defense (DoD), NSF, and NASA/NAS, have all adopted the TCP/IP suite as the most viable set of standards currently. As the international standards mature and products supporting them appear, it can be expected that the various networks will switch to using those standards.

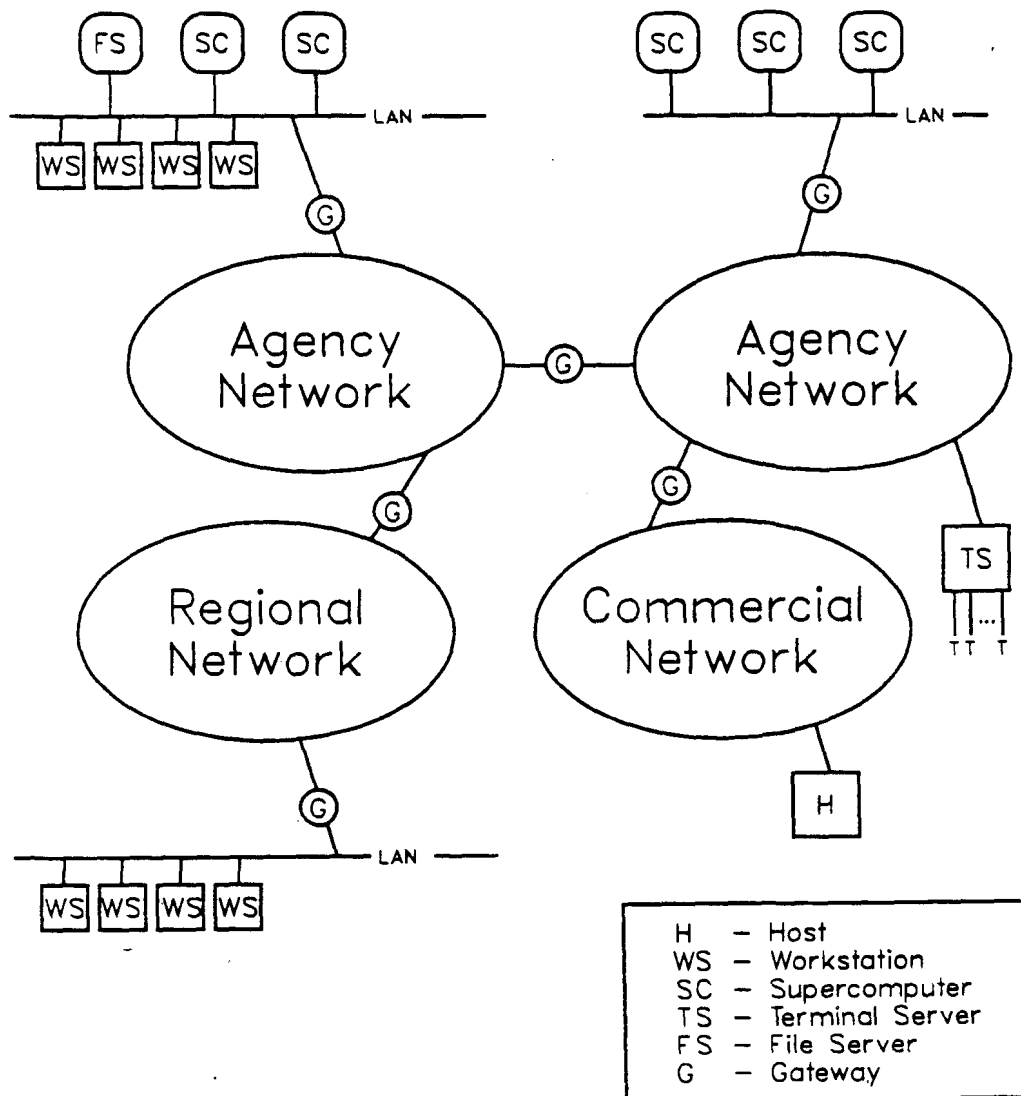


Figure 1. Internet system.

### 1.3. TECHNICAL APPROACH

The internet technology described above provides the basis for interconnection of the various agency networks. The means to interconnect must satisfy a number of constraints if it is to be viable in a multi-agency environment.

- a. Each agency must retain control of its own networks. Networks have been established to support agency-specific missions as well as general computer communications within the agency and its contractors. To assure that these missions continue to be supported appropriately, as well as assure appropriate accountability for the network operation, the mechanism for interconnection must not prevent the agencies from retaining control over their individual networks.\*
- b. Appropriate access control, privacy, and accounting mechanisms must be incorporated. This includes access control to data, resources, and the networks themselves, privacy of user data, and accounting mechanisms to support both cost allocation and cost auditing.<sup>23</sup>
- c. The technical and administrative approach must allow (indeed encourage) the incorporation of evolving technologies. In particular, the network must evolve towards provision of high bandwidth, type of service routing, and other advanced techniques to allow effective use of new computing technology in a distributed research environment.

#### 1.3.1. Communications Infrastructure

The communications infrastructure provides connectivity between user machines, workstations, and centralized resources such as supercomputers and database machines.\*\* There are two different types of networks. The first are local networks, meaning those that are internal to a facility, campus, etc. The second are networks that provide transit service between facilities. These transit networks can connect directly to computers, but are evolving in a direction of connecting local networks. The networks supported by the individual agencies directly are mainly in the category of transit (or long-haul) networks, as they typically provide nationwide connectivity, and usually leave communications within a facility to be dealt with by the facility itself. The IRI communications infrastructure thus deals mainly with the interconnection of transit networks.

The internet model described above provides a simple method for interconnecting transit networks (as well as local networks.) By using IP gateways between the agency networks, packet transport service can be provided between computers on any of the various networks. The placement of the gateways and their capacity will have to be determined by an initial engineering study. In addition, as the IRI evolves, it may be cost effective to install one or more wide area networks (or designate certain existing ones) to be IRI transit networks to be used by all agencies on a cost-sharing basis. Thus, the IRI communications infrastructure would consist of the interconnecting gateways plus any networks used specifically as transit networks. Using IP as the standard for interconnection of networks and global addressing provides a common virtual network packet transport service, upon which can be built various other network services such as file transfer and electronic mail. This will allow sharing of the communication facilities (channels, satellites, etc.) between the various user/agency communities in a cost-effective manner.

---

\* This is not to say that agencies may not choose to have their individual networks operated by the IRI, or even turned over to the IRI if they determine that to be appropriate.

\*\*This roughly corresponds to communications services at and below the transport layer in the ISO OSI reference model.

To assure widespread interconnectivity, it is important that standards be adopted for use in the IRI and the various computers connected to it. These standards need to cover not only the packet transport capability but also must address all the services required for networking in a scientific domain, including but not limited to file transfer, remote login, and electronic mail. Ultimately it is desirable to move towards a single set of standards for the various common services, and the logical choice for those standards are those being developed in the international commercial community (i.e., the ISO standards). However, many of the scientific networks today use one or more of a small number of different standards; in particular, the TCP/IP protocol suite mentioned above, the MFEnet protocols, and DECnet. As the international standards mature, it is expected that the number of communities using the same protocol suite will grow\* and therefore the ability of the users to share resources and results will increase.

### 1.3.2. User Services

So that scientists can effectively use the network, there needs to be a user support organization. To maximize the cost effectiveness of the overall IRI, the local user support personnel must be used effectively. In particular, it is anticipated that direct support of users/researchers would be provided by local support personnel. The IRI user support organization would provide support to those local support personnel in areas where nationwide common service is cost effective.

In particular, this organization would have several functions:

- a. assist the local support personnel in the installation of facilities compatible with the IRI;
- b. provide references to standard facilities (e.g., networking interfaces, mail software) to the local support personnel;
- c. answer questions that local personnel are not able to answer; and
- d. aid in the provision of specific user community services, e.g., database of relevance to specific scientific domain.

### 1.3.3. Internet Research Coordination

To evolve internet to satisfy new scientific requirements and make use of new technology, research is required in several areas. These include high-speed networking, type of service routing, new end-to-end protocols, and congestion control. The IRI organizational structure can assist in identifying areas of research where the various agencies have a common interest in supporting in order to evolve the network, and then assist in the coordination of that research.

## 1.4. MANAGEMENT APPROACH

A management approach is required that will allow each agency to retain control of its own networking assets while sharing certain resources with users sponsored by other agencies. To accomplish this, the following principles and constraints need to be followed.

- a. IRI consists of the infrastructure to connect agency networks and the user services required for effective use of the combined networks and resources.
- b. An organization must be identified to be responsible for the engineering, operation, and maintenance of both the interconnecting infrastructure and the user services support.

---

\*Even today, several of the agencies/communities are using a common protocol suite, namely the TCP/IP suite. All the users connected to those computers and networks are able to have the full functions of an interoperable networking capability.

- c. While some agencies may choose to make use of IRI facilities and contractors to manage their individual agency networks, this would not be required and is not anticipated to be the normal situation. Any such arrangement would have to be negotiated individually and directly between the agency and the IRI operations organization. Normally, the IRI organization would neither manage the individual agency networks nor have any jurisdiction within such networks.
- d. Gateways that interconnect the agency networks as well as any long-haul networks put in place specifically as jointly supported transit networks (if any such networks are required) will be managed and operated under the IRI organization.
- e. A support organization for common IRI services is required. The principal clients for these services would be the local support personnel.
- f. The IRI structure should support the coordination of the individual research activities required for evolution and enhancement of the IRI.

#### 1.4.1. General Management Structure

Figure 2 shows the basic management structure for the IRI. It is based on the use of a nonprofit organization (call it the Interagency Research Internet Organization, IRIO) to manage both the communications infrastructure and user support. The IRIO contracts for the engineering, development, operations, and maintenance of those services with various commercial and other organizations. It would be responsible for providing technical and administrative management of the contractors providing these functions. Having the IRI operational management provided by an independent nonprofit organization skilled in the area of computer networking will permit the flexibility required to deal with the evolving and changing demands of scientific networking in a cost-effective manner.

Direction and guidance for the IRIO will be provided by a Policy Board consisting of representatives from the government agencies who are funding the IRI. The Chairman of the Board will be selected from the agency representatives on a rotating basis. The Board will also have an Executive Director to provide administrative and other support. To provide effective support for the IRI Policy Board as well as assure appropriate coordination with the IRIO, the Executive Director shall be the Director of the IRIO.

To assure that the IRI provides the best support possible to the scientific research community, the Policy Board will be advised by a Technical Advisory Board (TAB) consisting of representatives from the network research and engineering community, the various networks being interconnected with the IRI, and the scientific user community. Members of the TAB will be selected by the Policy Board. The TAB will review the operational support of science being provided by the IRI and suggest directions for improvement. The TAB will interface directly with the IRIO to review the operational status and plans for the future, and recommend to the Policy Board any changes in priorities or directions.

Research activities related to the use and evolution of the internet system will be coordinated by the Internet Research Activities Board (IRAB). The IRAB consists of the chairmen of the research task forces (see below) and has as ex-officio members technical representatives from the funding agencies and the IRIO. The charter of the IRAB is to identify required directions for research to improve the IRI, and recommend such directions to the funding agencies. In addition, the IRAB will continually review ongoing research activities and identify how they can be exploited to improve the IRI.

The Research Task Forces will each be concerned with a particular area/emphasis of research (e.g., end-to-end protocols, gateway architectures, etc.). Members will be active researchers in the field and the chairman will be an expert in the area with a broad understanding of research both in that area and

the general internet (and its use for scientific research). The chairmen of the task forces will be selected by the IRAB, and thus the IRAB will be a self-elected and governing organization representing the networking research community. The chairmen will solicit the members of the task force as volunteers.

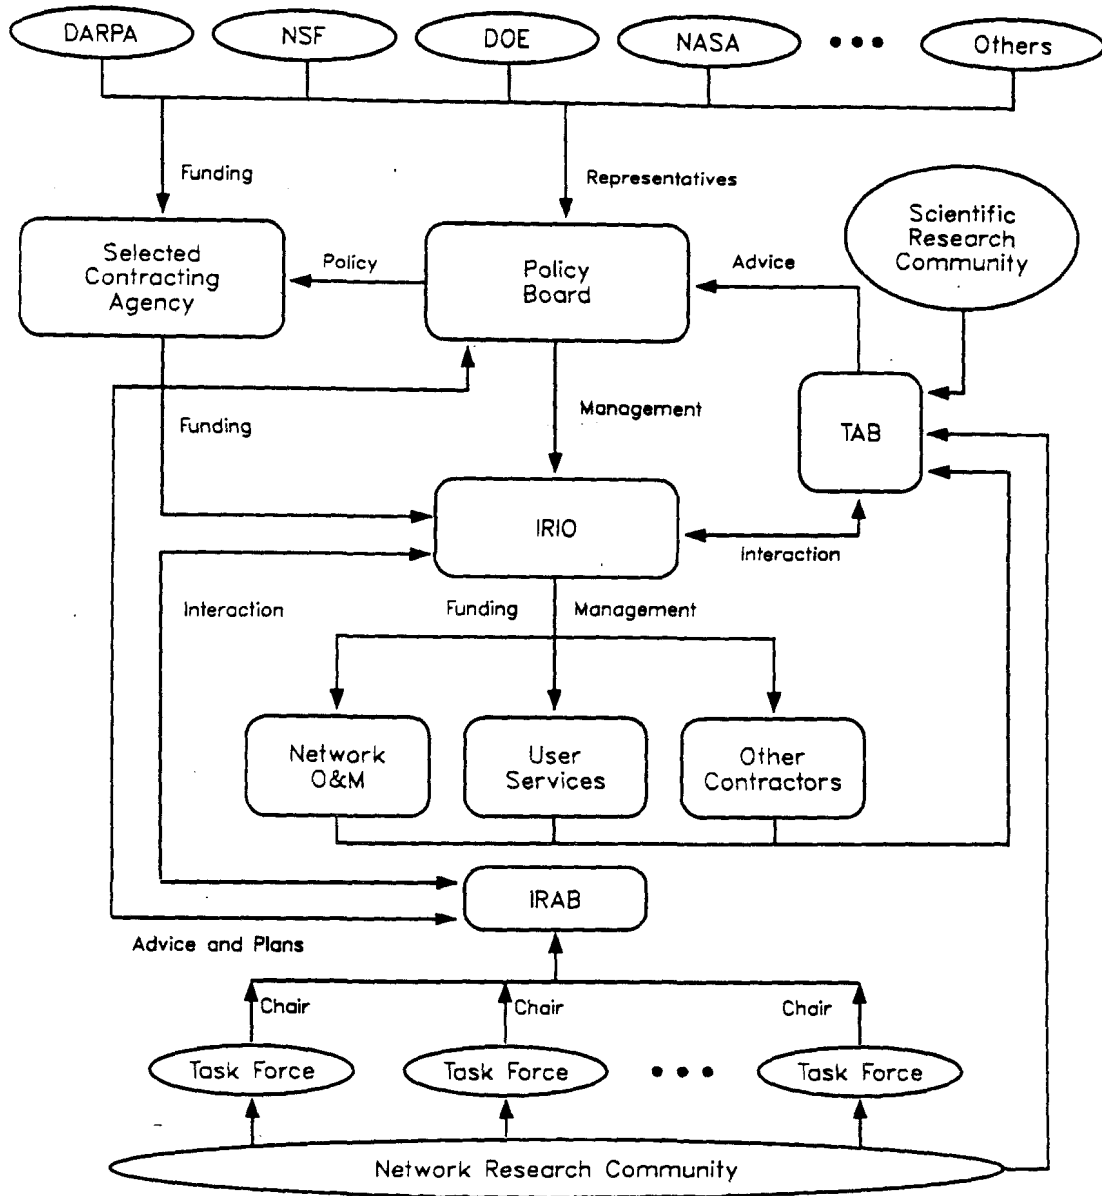


Figure 2. IRI management structure.



## 1.4.2. Funding

In this section, the funding of the IRI is described. Recall that the IRI consists of the infrastructure to connect the agency networks and the services required for users to make effective use of such an infrastructure. These costs are divided into two categories: operations costs and research costs. The operations costs are those to operate and maintain both the communications infrastructure and the user services. These costs must be shared between the various agencies and channeled to the IRIO to operate the IRI. The research costs are those used to carry out the needed research to evolve the IRI. These costs are handled within the various agency budgets and used to support research in each agency with coordination between the agencies.

### 1.4.2.1. Operations Costs

Each participating agency will contribute a share of operations cost of IRI. Initially, each agency will contribute an equal share. Later, perhaps, the agency contributions will be adjusted according to a number of factors such as number of users, amount of traffic, and type of support required (high bandwidth real time versus low bandwidth mail, for example).

To facilitate the funding and administration of the IRI, one agency will be selected to manage the contract with IRIO. All funds will flow through that agency to the IRIO via interagency transfer. The role of the selected agency would be to provide the needed contractual activities and administrative management. Technical guidance and monitoring of IRIO activities would be provided by the IRI Policy Board.

It is not yet clear which federal agency is best for this role. The requirements for such an agency include the ability to deal flexibly with the evolving requirements of the IRI, to deal with funding flowing from the various agencies, and to deal flexibly with the various agency technical representatives and incorporate their recommendations into the contract as required. One of the first activities required for the Policy Board would be to select an appropriate funding agency.

All operations and maintenance funding for the IRI will flow through the IRIO to selected contractors. This allows centralized management of the operation of the IRI.

There are two major assumptions underlying the budgetary estimates. First of all, the IRIO should maintain a fairly low profile with respect to the end users (i.e., the scientists and researchers). That is, the users will interact directly with their local support personnel. The IRIO will act as facilitator and coordinator, and provide facilities, information, and help services to the local sites. This will allow the IRIO to remain relatively small, as it will not need to deal directly with the thousands of scientists/users. Second, it is assumed that the operations budget supports the interconnection of agency networks as well as transit networking where required, but does not include costs of the individual agency networks.

Section 1.7 provides details of the budgetary estimate and Table 1 gives a summary. Note that the initial year has a higher expenditure of capital equipment, reflecting the need to purchase both the gateways needed for initial interconnection and the needed facilities to provide the operation of the gateways and the user services. Operations costs are expected to grow by inflation while the capital costs should remain constant (decrease when inflation is considered) as the IRI is stabilized.

Fiscal Year	Capital Costs (\$M)	O&M Costs (\$M)	Total (\$M)
1987	2	8	10
1988	1	9	10
1989	1	10	11
1990	1	11	12
1991	1	12	13

### 1.4.2.2. Research Costs

In addition to the costs of operating and maintaining the communications infrastructure and user services, funding must be allocated to support an ongoing program of research to improve and evolve the IRI.

While each agency funds its own research program, the intent is that the various programs are coordinated through the IRI Policy Board. Likewise, while it is not intended that funds shall be combined or joint funding of projects is required, such joint activity can be done on an individual arrangement basis.

Each agency agrees, as part of the joint IRI activity, to fund an appropriate level of networking research in areas applicable to IRI evolution. The total funding required is currently estimated to be four million dollars in FY87, growing by inflation in the outyears. Details of this budgetary estimate are provided in Section 1.7.

## 1.5. PHASED IMPLEMENTATION PLAN

The long-term goal of the IRI activity is to put in place a functional high-performance network available to scientists across the nation. To accomplish this goal, a steady evolution of capability is envisioned. This phased approach involves both technical and administrative aspects.

### 1.5.1. Technical Phasing

Currently, networks are being supported by a number of agencies as discussed in Section 2. Many are using the DoD protocol suite (TCP/IP, etc.) and others have incorporated or are incorporating mechanisms for interoperability with networks using the DoD protocol suite (e.g., MFEnet). Most have discussed eventual evolution to ISO protocols and beyond. By and large, most of these networks are hooked together in some mainly ad hoc manner already, some by pairwise arrangement and some through third party connections (e.g., a university network connected to two agency networks).

There are two major shortcomings to this ad hoc connection, though. Performance is not adequate for advanced scientific environments, such as supercomputer usage, and community-wide user support is not generally available. The phased approach described below will allow these deficiencies to be overcome through coordinated action on the part of the various funding agencies.

### 1.5.1.1: Phase I--Functional Interoperability

The initial stage of the IRI would provide for sharing of the communications facilities (e.g., channels, satellites, etc.) by interconnecting the networks using the IP and IP gateways. In addition, mechanisms will be installed (where required) and maintained to allow interconnection of the common user services, such as electronic mail. This will allow sharing of resources attached to the network, such as supercomputers.\*

Specific steps to be undertaken in Phase I are the following:

- a. Gateways will be purchased and installed where needed to interconnect the agency networks. The location and performance of these gateways will be specified by the IRIO and approved by the Policy Board. This engineering will take into account an estimate of current and future traffic requirements as well as existing interconnecting gateways. It may also result in a recommendation that some or all existing gateways between agency networks be replaced with common hardware so that adequate management of the interconnection can be achieved.
- b. An IRI operations and management center will be established for the interconnecting gateways.\*\*
- c. The requirement for application gateways or other techniques to interconnect communities using different protocols will be investigated and a recommendation made by the IRIO in conjunction with the IRAB. The appropriate mechanisms will be installed by the IRIO at the direction of the Policy Board.
- d. An initial user services facility will be established. This facility will provide at a minimum such services as a white pages of users (similar to the current internet "whois" service) and a means for making accessible standard networking software.
- e. The IRAB, in coordination with the Policy Board, will draft a coordinated research plan for the development of the new technologies required for evolution of the IRI.

### 1.5.1.2. Phase II--Full IRI Capability

Phase II will make the IRI fully functional with enhanced capabilities and performance.

- a. High-performance gateways with appropriate new capabilities and functions will be installed, replacing and/or augmenting the gateways in place from Phase I. The functionality and performance of these gateways will be specified based on the experience from Phase I use, the anticipated new uses of the network, and the state-of-the-art technologies available as a result of the ongoing research.
- b. The basic user services facility will be mature and support network operation. New capabilities will be developed to support specific scientific communities (such as a database of software used by a specific community and its availability over the network).
- c. A high-performance backbone network will be installed if needed to connect high-performance agency networks.\*\*\*

---

\*Note: actual use of facilities other than mail would require arrangements with the various responsible parties for each host. For example, to log in to a host not only requires network access; it also requires a login account on that host.

\*\*This perhaps could be done in conjunction with a network management center for another set of gateways, e.g., those supported by DARPA or NSF.

\*\*\*This is anticipated because of the move in several agencies to provide high bandwidth networks in support of such activities as super-computer access.

- d. The introduction and use of international standards will be investigated and a plan developed for providing more services to the broad scientific community through use of these standards.

### 1.5.2. Administrative Phasing

The goal of the IRI is to get to a fully cooperating and managed interagency research internet involving most if not all of the agencies supporting scientific research. Recognizing that currently the major research networking players (both networking for research and research in networking) are DOE, NASA, DARPA, and NSF, the following steps are recommended.

- a. The first and critical step is to establish a four agency Memorandum of Agreement to interconnect the agency networks and to share the costs of interconnection, transit networks, and an operations center. A management structure should be agreed upon as outlined above. Agreement must also be reached on the need to fund an ongoing research and engineering activity to evolve the internet.
- b. A Policy Board and TAB should be established as quickly as possible to assure appropriate guidance and direction.
- c. The Policy Board shall then select an agency to handle the administrative and contractual actions with the IRIO.
- d. A nonprofit organization shall then be selected by that agency through an appropriate procurement mechanism to be the IRIO. The Policy Board of the IRI shall be the selection panel.
- e. The initial four agencies shall transfer the agreed upon funds to the selected contracting agency on equal basis to start.
- f. These funds will then allow the contracting agency to establish a contract for the IRIO with the selected nonprofit organization.
- g. The IRIO can then establish subcontracts for engineering, procurement, installation, and management of gateways and operation of the user services center.

To initiate the research coordination, the following steps will be accomplished.

- a. The Internet Activities Board will evolve into the IRAB through added membership and charter revision.
- b. Additional task forces will be formed as needed to reflect the expanded areas of research interest.

Once the IRI is established and operating, the funding and use of the IRI will be reviewed to determine if equal funding is equitable. If not, the IRIO should be tasked to develop a recommendation for a practical cost allocation scheme. In addition, once the IRI has proved itself to be successful, other agencies will join the IRI and provide additional funding.

### 1.6. INDUSTRY ROLE

This paper has thus far addressed the interconnection of agency supported networks and the use of such an internet by agency supported researchers. However, industry also has a need for a similar infrastructure to support its research activities.\* Regulatory concerns make it difficult for industry to connect to a network that is supported by a federal agency in pursuit of the agency mission.

---

\*Note that this refers only to industrial research activities. It is not envisioned, nor would it be appropriate, for the IRI to provide a communications system for normal industrial activities.

The IRI structure above, though, may permit the connection of industrial research organizations. Since the IRIO is a nonprofit, nongovernment organization, it would be able to accept funds from industry as a fair share of the costs of using the IRI. These funds in turn can be used to expand the networking resources so that no degradation of service is felt by the users supported by the federal agencies. This topic would need to be discussed further by the Policy Board and the organization selected as the IRIO.

## 1.7. FUNDING BREAKDOWN

This section provides the details for the budgetary estimates of Table 1.

### a. Gateways

Gateways will be required between the various agency (and perhaps regional) networks. As an upper bound, assume one IRI gateway per state times \$40K per gateway, spread out over 2 years, for a capital cost of \$1M per year for first 2 years.

### b. Operation Center

The IRI operations center will have to engineer the location and capacity of the gateways, as well as install, operate, and maintain them. It also will need to coordinate support and maintenance of end-to-end service, helping to identify and correct problems in the interconnections. Costs are estimated as two people round the clock to man the operations center and three full time people to coordinate, operate, and engineer the IRI. Using an estimate of \$120K [including other direct costs (ODC)] per year for an operator and \$200K per year for other activities, and translating two people round the clock into nine people results in a total annual cost of \$1.7M. In addition, equipment costs of roughly \$500K per year can be expected.

### c. Transit Networks

It is expected that support of at least one transit network will be necessary. This may involve reimbursement to one of the agencies for use of their network, or may involve operations and maintenance of an IRI dedicated network. An estimate for these costs, based on historical data for operating the ARPANET, is \$4M per year.

### d. User Support Organization

To provide effective support as discussed above will require a staff available during working hours. A reasonable estimate for the costs of such an organization is five people times \$200K per year, or \$1M per year (including ODC). In addition, there will be capital equipment costs in the first 2 years totaling roughly \$2M.

## 1.8. SUMMARY AND CONCLUSIONS

The interconnection of the various agency networks supporting scientific research into an overall infrastructure in support of such research represents an exciting opportunity. This report recommends an approach and a specific set of actions that can achieve that goal. It is hoped that, regardless of the mechanism used, that the federal agencies involved recognize the importance of providing an appropriate national infrastructure in support of scientific research and take action to make such an infrastructure a reality.

## 1.9. ACKNOWLEDGMENT

This paper was prepared with advice and comments from a large number of people, including the members of the FCCSET Committee Network Study Group and the Internet Activities Board. Their input is greatly appreciated, and I hope that this report represents a consensus on both the need for the IRI and the proposed approach.

## 1.10. REFERENCES

1. *Report on Interagency Networking for Research Programs* (February 1986).
2. V. G. Cerf and P. Kirstein, "Issues in Packet-Network Interconnection," *Proceedings of the IEEE*, pp. 1386-1408 (November 1978).
3. V. G. Cerf and E. Cain, "The DoD Internet Architecture Model," *Computer Networks*, pp. 307-318 (July 1983).
4. B. M. Leiner, J. Postel, R. Cole, and D. Mills, "The DARPA Internet Protocol Suite," *IEEE Communications Magazine* (March 1985).
5. *A History of the ARPANET: The First Decade*, Defense Advanced Research Projects Agency, Defense Technology Information Center report AD A1 15440 (April 1981).
6. I. M. Jacobs, et al., "General Purpose Satellite Networks," *Proceedings of the IEEE*, pp. 1448-1467 (November 1978).
7. F. Tobagi, R. Binder, and B. M. Leiner, "Packet Radio and Satellite Networks," *IEEE Communications Magazine* (November 1984).
8. P. E. Kahn, et al., "Advances in Packet Radio Technology," *Proceedings of the IEEE*, pp. 1468-1496 (November 1978).
9. D. Clark, et al., "An Introduction to Local Area Networks," *Proceedings of the IEEE* (November 1978).
10. J. Lederberg, "Digital Communications and the Conduct of Science: the New Literacy," *Proceedings of the IEEE*, 66, pp. 1314-1319 (November 1978).
11. J. C. Hoskins and J. S. Quarterman, "Notable Computer Networks," *Communications of the ACM* (October 1986).
12. P. J. Denning, A. C. Hearn, and C. W. Kern, "History and Overview of CSNET," *Proc. ACM Sigcomm Symposium*, pp. 135-145 (March 1983).
13. D. Comer, "The Computer Science Research Network CSNET: A History and Status Report," *Communications of the ACM* 26, pp. 747-753 (October 1983).
14. F. R. Bailey, "NAS: Supercomputing Master Tool for Aeronautics," *Aerospace America*, 23, pp. 118-121 (January 1955).
15. D. M. Jennings, L. H. Landweber, I. H. Fuchs, D. J. Farber, and W. R. Adrion, "Computer Networking for Scientists," *Science*, 231, pp. 943-950 (February 1986).
16. V. G. Cerf and R. E. Kahn, "A Protocol for Packet Network Intercommunication," *IEEE Transactions on Communications* COM-22 (May 1974).
17. H. Zimmerman, "OSI Reference Model - the ISO Model of Architecture for Open Systems Interconnection," *IEEE Transactions on Communications* COM-28, pp. 425-432 (April 1980).
18. *MIL STD 1777: Internet Protocol*, Defense Communications Agency (1983).

19. *MIL STD 1778: Transmission Control Protocol*, Defense Communications Agency (1983).
20. *MIL STD 1780: File Transfer Protocol*, Defense Communications Agency (1985).
21. *MIL STD 1781: Simple Mail Transfer Protocol*, Defense Communications Agency (1985).
22. *MIL STD 1782: Telnet Protocol*, Defense Communications Agency (1985).
23. B. M. Leiner and M. Bishop, "Access Control and Privacy in Large Distributed Systems", Research Institute for Advanced Computer Science report RIACS TR 86.6 (March 1986).





## 2. THE NASA SCIENCE INTERNET

---

**Wm. Prichard Jones**  
NASA Ames Research Center

The National Aeronautics and Space Administration (NASA) Science Internet (NSI) is a full-service, packet switched network that provides remote login, electronic mail, and file transfer services utilizing standard, vendor-independent internet and transmission control protocols. The NSI will link electronically the science communities together, including the disciplines of astrophysics and astronomy, as well land, ocean, climate, and planetary science. These communities will also be tied to other interoperable networks such as the Defense Data Network (ARPANET, MILNET, MINET, etc.) and the National Science Foundation network (NSFnet), and be accessible also from public packet networks.

The Program Support Communications Network that NASA has implemented recently is the backbone of NSI. Sixteen NASA centers and facilities are currently connected with high-speed digital switches and earth stations at many of the locations. Terrestrial and satellite 1544-kbit/s backbone links carry both circuit switched and packet switched data as well as voice, facsimile, voice conferencing, and eventually, video conferencing services. Using a high-performance packet processor as a gateway at several geographically distributed regional hubs and additional packet processors at science sites, the agency can provide a cost-effective solution to the growing cost of point-to-point circuits by installing the NSI capability. Such a capability also can be more reliable than individual point-to-point circuits because redundancy in the backbone is affordable and tail circuits can be made redundant with a dial back-up capability.

A typical science site will operate a local area network such as an Ethernet or Token Ring to which a gateway processor is attached. This gateway is connected to the nearest regional hub gateway. One or more of the regional hubs will operate gateways to other networks. The current protocol suites that have been approved for use on the NSI are the Internet Protocol (IP) and the Transmission Control Protocol (TCP). With IP/TCP protocols comes a suite of application protocols to permit remote login, standardized mail, file transfer, and internet routing. A central clearinghouse to consolidate science community connection requirements and provide standardized hardware and software, configuration parameters, addressing, and user support is to be located at Ames Research Center, Moffett Field, California. This center has pioneered the effective use of computer communications for science programs.



### 3. THE ROLE OF REGIONAL NETWORKS

---

**Richard Mandelbaum**  
University of Rochester

The current model of the National Science Foundation network has a three level-hierarchy consisting of

- (1) the transcontinental backbone network,
- (2) the regional and consortia networks, and
- (3) the campus networks.

In order to discuss the role of Level 2 in this hierarchy, I will begin by making some assumptions about the responsibilities of Level 1 management.

- **Standards**--Level 1 management should set minimal hardware and software standards for both the backbone and Level 2 networks. These should include such things as media exchange standards and minimal levels of services standards as well as standards for network engineering and design.
- **Operations and Engineering**--Level 1 management should play the lead role in determining Level 1 coordinating and Level 2 network configurations and in monitoring operations. This should be done either by direct operation of appropriate network operation centers (NOCs) or by coordinating NOCs operated by regionals or groups of regionals. In addition, Level 1 would be responsible for global name and addressing specifications and coordination of the development and testing of future protocol modifications and additions.

Given the above, what is the primary obligation of the regionals?

Clearly, if Level 1 is setting standards, then Level 2 must implement them. The regionals must do most of the actual network topology and connection engineering and operations and maintenance. More crucially, the Level 2 networks must be responsible for information exchange. This would include technical and policy information on

- connecting to the network,
- hardware and software standards,
- local implementations and relationships to campus local area networks, and
- supercomputers and other network resources.

The regionals should develop user's lists and maintain telephone hot lines and computer bulletin boards. They should develop and maintain lists of contacts, up-to-date network maps, databases of network resources, and user interests (e.g., network "White Pages" and "Yellow Pages"). They should maintain and coordinate "networking circuit riders" who would assist local campuses with implementation of local area network connections and troubleshooting.

In addition, the regionals should sponsor seminars and workshops for both their campus-level technical contacts and campus end users, as well as produce newsletters and other information bulletins.

They should also service the needs of discipline-oriented user groups (e.g., High Energy Physics, Computer Science, and medicine and library groups such as RLG or OCLC) and provide discipline-based network assistance. In addition, the regionals should support a certain amount of "experimental" network activity by providing support for testing of protocol modifications, new routing algorithms, etc.

On a political level, the regionals should be the primary contacts with state governments and local industries, especially the telephone companies. They should develop regional based support for networking by

- expanding the range of user's who need a computer network, and
- using networking to help expand the breadth of academic industrial contacts.

## NETWORKING TRENDS



## 4. NETWORKING: SOME OBSERVATIONS ON THE FUTURE

---

Jack Haverty  
BBN Communications Corporation

### Abstract

*The following discussion presents several opportunities and observations on the future evolution of a communications infrastructure. This paper was originally prepared in 1986 as part of an effort to address the future use of communications in government applications. I've modified this discussion to address the context of a communications infrastructure to support national research activities.*

*The examples that follow focus on cases where most of the required technology has been demonstrated and the basic feasibility/research cycle has occurred, but the techniques or approaches are not yet in general use.*

### 4.1. INTRODUCTION

During the past decade, there has been an explosion in communications and computer technologies. In communications, a wide repertoire of technologies are now on the menu, each with associated costs, capabilities, and availability. In the computer area, the expansion of available systems has been dramatic, ranging from nanoproducts to supercomputers.

In both communications and computers, vast shifts in costs-per-capability have occurred, with resultant shifts in system architectures. The effects of these shifts on fielded systems are just beginning to be noticed with the introduction of elements such as widely distributed high-performance workstations.

In general, system structures are tending toward a new architecture. Previous systems utilized a small number of high-performance processing centers, with communications to a much larger number of unintelligent devices such as sensors and terminals. New systems will utilize much more computing capacity that will be widely distributed, with a smaller number of very powerful computing facilities in the communications web.

Architectural shifts have resulted in a growing spectrum of communications requirements. Where communications in past years might have been divided into a few classes, such as interactive and bulk, communications in the future will need to serve a wider range of requirements, such as supporting voice, data, graphics, and other uses. Low-delay, low-bandwidth control traffic, interactive, transaction, bulk, multidestination, and message traffic, both for human and computer-computer communications, will also be needed.

In addition to a growing spectrum of communications technologies, there is a growing trend toward technologies that provide on-demand services. Shared satellite channels, dialup circuits, controllable multiplexing devices, and other such systems offer a means of managing a communications system in real time, that is, in response to the user's immediate needs. Because such technologies are relatively expensive, high-performance communications facilities will need to be accessible to a large number of users on a shared basis, much as early computer facilities were shared through batch and later through timesharing techniques. The opportunity here lies in applying systems approaches to creating

communications systems that can effectively use a repertoire of communications technologies to meet a range of communications requirements in the most cost-effective way consistent with the situation at the time and the policy for allocating resources to meet user requirements.

In a sense, the thrust of this idea is to view the communications infrastructure in the context of a management task. This approach should be applied at two levels. The underlying transmission systems can be viewed and managed as a set of assets that can be applied as required to serve a variety of users; the diverse user requirements can be viewed as a set of requests for different types of service. The management task is to make the most effective decisions and scheduling to meet the requests of the community, and to implement those decisions by control of the communications facilities.

A satellite system that carries voice in the morning might be used to transport bulk data during the night. A message system might normally utilize a high-bandwidth transoceanic trunk, but it could be temporarily interrupted to provide communication support for a time-critical database transfer.

The ability to integrate these various technologies and to support a variety of user applications sharing the same assets can have profound effects not only on the costs of a system operation, but also on the overall availability. The goal is to be able to utilize any available communications asset to support the highest priority usage at the time. One approach to this goal would be to apply the packet switching paradigm to multidisciplinary communications system; i.e., constantly measure the system capabilities, measure mission service demands, and then decide how to apply resources (i.e., choose technology for particular communications requests) to best meet demand with available resources in a distributed system.

## 4.2. NETWORKS AS COMPUTERS

Traditional approaches to the use of computers and communications have involved "running" a program on a computer, allowing it access to communications to move data and to interface input and output devices such as printers or humans at terminals. The communications technology in this case performs a pure transport function.

Relatively recent introductions to the technology of data communications have included various kinds of "servers" that perform a variety of functions such as protocol conversion, printing, database storage and retrieval, file storage, and the like. In the marketplace, the local area network and workstation arena has been the focus of these activities.

This technology is an instance of a more fundamental technology that is now applicable in the overall computer/communications arena, and has the potential for a profound impact on the capability and effectiveness of widely distributed systems such as are becoming more and more pervasive.

To illustrate this technology, consider the analogy of a large-scale data processing installation with a large-scale, geographically widespread, distributed system such as a typical current database application. In such a comparison, the communications substrate, containing wires, satellites, fiber, radio, and other transmission systems, is analogous to the backplane of a traditional computer system. Both perform the function of transporting information, without significant processing of any kind relevant to the overall user application.

To make a data processing installation effective to its users, other elements are of course integrated into the system. Memory, mass storage, output devices such as printers or microfilm, input devices



such as tapes and sensors, and communications devices to interface to users or other computing facilities are all configured, enhanced, and operated to make the overall system meet the needs of its client. Software is procured or written and modified as needed to maintain effectiveness in a changing environment.

Introductions to communications technology such as the various servers mentioned above are analogous to the introduction and standardization of elements of traditional distributed processor (DP) installations. Over time, the architecture of computers has evolved from the Von Neumann stored-program approach to include the finer details of multiprocessors, specialized processors such as data channels, various types of mass storage with different capabilities, and associated costs, etc. In building a DP installation, methodologies now exist for evaluating alternatives and making decisions about system configuration, with of course the associated and available repertoire of techniques, products, and services.

The key aspect of the data processing paradigm is that the overall end-user requirement is being addressed by putting together an integrated system. The opportunity in this area is to exploit the idea of extending the concept of a wide area communications system to include functions other than simple transport.

In the local area technology, this has already begun with development of various kinds of servers. Single-vendor approaches have been able to create integrated system environments of servers, workstations, and workhorse processors. Such integrated environments provide high performance and capability compared to interconnections of components from multiple vendors. In computer system technology, many procedures, concepts, and architectures have evolved over the last two decades. Notions such as caches, I/O subprocessors, data storage methods and access techniques, operational procedures, and capacity planning have all evolved as computer technology itself has matured.

Even in the wide area technology, these kinds of trends are beginning to be evident. For example, several current networking products have mechanisms whereby a user must "log in" to the network, much as he would to a timesharing computer.

One approach to capitalizing on these technologies is to pursue a systems approach to the overall communications infrastructure by introducing technologies from computer technology to obtain advantages from that analogy and expanding the concepts now appearing in local environments to include more geographical spread. The concept of viewing a communications network as analogous to a computer system leads to the idea of moving processes rather than data (see below); this is simply the notion of programming the "computer" extended to a distributed environment.

#### **4.3. MOVING PROCESSES INSTEAD OF DATA**

The abstract model of computation includes three basic elements: the input, the processing itself, and the output. In a particular application that utilizes computation, these elements are often arranged in a feedback loop, where the inputs are derived, or sensed, from the real world, and the outputs affect the real world, either directly or as a result of some human decision-making based on observation of those outputs.

Current approaches to applying computers and related technology to real-world problems almost universally involve systems wherein data are moved over communications media, from processor to processor, in order to perform some task. For example, sensor data are detected and often processed at the sensor to generate some data, which in turn are communicated to an analysis site, which further

aggregates and processes data. Then this structure is repeated, often in a hierarchical way. The task of constructing a particular system involves defining the particular processing to be done in at each site and the communications requirements between the various sites in terms of bandwidth, delay, and topology.

This structure is probably a natural result that translates the techniques used in human daily life activities into the world of computing. The development of communications in the human world has enabled the creation of such mechanisms as newspapers, radio, television, letters, telegrams, and other such communications mechanisms, each of which has different characteristics. Human interaction then develops by applying appropriate technologies to each requirement. For example, television is used to move summary data of current events quickly to large numbers of people. Newspapers move more data, although less quickly to a broad audience. Journals move in-depth data to a fairly narrow audience and often carry data referring to events that occurred weeks or years earlier. All of these cases involve moving data to the processing elements, i.e., the humans who will use that data to make decisions about their daily lives.

Other mechanisms also exist in human activity, however. For example, libraries are repositories of information that needs less wide, less immediate distribution, or for which the sheer volume of information precludes wide distribution due to cost considerations. People can borrow material from libraries, and mechanisms such as bookmobiles exist, but the use of libraries primarily involves people going to the library and processing the data there, and leaving with increased knowledge or notes representing the results of their processing of the data.

Computer communications have primarily been oriented toward the development of systems that fit the first set of mechanisms, namely moving the data to the processors. In fact, more and more technologies have been created for performing the transport functions, with different costs and capabilities, much as different publication or transportation systems have evolved in the arena of human activities.

In this architecture, an initial system design dictates the locations of the processing functions for the lifetime of the system. Data are transported between sources, processing elements, and sinks to carry out the function of the overall system.

As a system is used over time, things change. When it is determined that a processor has insufficient power, or inadequate instructions about what to do, a component subsystem must be changed by replacing the hardware or software as necessary. This is a relatively slow process in current fielded systems.

The opportunity in this area is to exploit the other model for merging computers and communications to perform any particular function, namely moving the processing to the data. In human terms, this has traditionally meant moving people, either physically or by proxy. For example, much data exist at the library, but I might either go there myself or send an assistant with instructions about the processing I desire to have done.

In the computer information processing world, it is not necessary to move the computers, but it is feasible to move the processes. One can envision a typical current system in which data are collected, kept in repositories, and processed to generate new data, which are also stored in repositories and sent to other processing centers over communications lines. Most current systems have this behavior.

However, in addition, one can imagine new methodologies in this processing environment, in which processes likewise move, using the communications systems as transportation accompanied by luggage

representing the data that have been collected. A user in such a world would build a process at his home site and launch it into a "network" of communications media, machines, and databases, with explicit instructions about things to look for and what to do when they are found.

This methodology is especially promising in situations where quick response to an event is required. Some work has already been done in this area, in particular by Schoch at Xerox Parc, where a system was developed enabling a process called a "worm" to travel from machine to machine in a local network environment. Early ARPANET experiments in the 1970s similarly explored these concepts and performed some simple demonstrations.

The availability of artificial intelligence technology also makes this kind of approach more and more applicable, since it is possible to create autonomous programs that can function without as much need for frequent human interaction as in traditional systems.



## 5. FUTURE DIRECTIONS IN COMMUNICATIONS RESEARCH

---

David D. Clark  
Massachusetts Institute of Technology

### Abstract

*Over the last two decades, communications networks have transformed the basic nature of computing by connecting individual computers into a vast mesh of information and resources. Unforeseen computing developments, such as remote login and resource sharing, have become commonplace to facilitate the interconnection. Even now, we cannot assess the final impact that networking will have on computing, because networking represents a fundamental extension to the base of computing, and we have not yet had long enough to understand how best to build on this base. Electronic mail, for example, is an entirely new and unpredicted application that has extended and changed the way humans deal with each other. This exploration of computer-mediated human communication has in turn suggested new ways for computers to communicate. In this way, our understanding of computer communications continues to grow.*

*An early technological innovation that has assisted in computer networking has been packet switching. During the last 20 years, while packet switching has grown, additional new and exciting technology innovations have occurred--most importantly, very large scale integration (VLSI) and fiber optics. VLSI and fiber optics are permitting us to build a new generation of vastly powerful communications networks. We must now decide what capabilities computing networks ought to have, although we will probably fail to predict the exact eventual impact that these new networks will have.*

### 5.1. TECHNOLOGY

Fiber optics offer the possibility of vastly increased bandwidth at vastly reduced costs. The fastest link commonly found in the wide area networks today is 10 Mbit/s, while 100-Mbit local area networks are now available and 1-Gbit links are possible. Over the next 10 years, further increases of 10 or 100 Gbit/s can be expected. These speeds, if exploited, would completely change the capability provided by the communications substrate. One could imagine a nationwide network of computers communicating at speeds previously undreamed of. Equally important, such a network could completely change the economics of communication by removing bandwidth as the principal cost component of networks. This change would be sweeping; it would do for communications what the inexpensive microprocessor chips did for computing.

Achieving high bandwidth is not just a matter of installing new high-speed fibers. The current communication bottlenecks are not the links but the switching nodes and the host interfaces. Currently, the most sophisticated switching nodes are hard pressed to keep up with a 10-Mbit/s link, and they will require a fundamental redesign if they are to match the speeds that fiber will bring. To take advantage of the speeds that will come in the next decade, we will almost certainly need to develop a new switching architecture, for neither packet switching nor circuit switching seems to have the scope to meet the new demands.

The most important technology for a new network architecture will be very large scale integration (VLSI), because VLSI will permit the economic creation of new sorts of switches. Highly parallel switches could permit the forwarding of vast flows of data, and highly intelligent host interfaces could eliminate the performance barriers now found in protocol software. VLSI and fiber are thus two critical ingredients that combine to permit a leap forward in performance.

VLSI is critical in other ways as well. For example, most current switching nodes have rather small memories, since they are simply minicomputers, and the management algorithms for the network have been designed to match this constraint. Routing algorithms for today's packet switching networks, for example, are limited in capability, and could not possibly scale to permit a packet network the size of the current phone system. But with the memory costs of today, switching nodes could store and manage routing tables of a size unimagined 10 years ago. With VLSI to support a new generation of algorithms in the switching nodes, we can thus envision that our new generation of communications substrate can grow to be as universal as the phone system. All our experience in communications, both with the phone system and with data communications, tells us that universal access is a fundamental and desirable capability. VLSI is thus a critical capability for two reasons: it permits high speed and large size.

### 5.1.1. The Future of Computing Networks

As we have discussed above, we have the opportunity to build a new generation of computer communication networks. Indeed, we must address the future requirements that computing will have when this network is developed. In this section, we will examine several evolving aspects of computing in order to assess their potential impacts on computer communications.

### 5.1.2. Supercomputers

While the microprocessor has permitted the distribution of most computing power to the site of the user, the supercomputer will continue to be a centralized resource. The power of the current supercomputer may well appear in distributed form, but new and more powerful machines will certainly arise, so there will always be the need to interact with a remote processor. This realization is the basis for the current construction of networks in support of scientific computing.

Currently, the goal of supercomputer networks is to make the use of supercomputers interactive. If the programmer can see the output of the program while it runs, perhaps in graphic form, a faulty execution can be terminated, and real-time interaction with the program is possible. Real-time graphic interaction might require 1 Mbit/s, a rate that exceeds current network capabilities but that could be easily engineered.

More interesting is the speculation that in the future there will be multiple sorts of supercomputers that will work jointly on a problem. For example, in VLSI design we see numerical, symbolic, and graphic specialization within the design process. Although it is not clear what sorts of communication capability will be needed to tie these machine processes together, one can speculate that the bandwidth will be very high and the latency very low.

### 5.1.3. Parallel Computing

A currently popular view is that multiple computers working in parallel can be an effective and economical source of massive computing power. These sorts of configurations have massive communications requirements. The requirements are rather specialized--very short paths but very low

delays--and specialized networks are usually proposed to meet them. However, it is not clear whether this specialization should exclude this application from the scope of general communications facilities, since new switching facilities may well have the correct latency and speed capabilities to support this application.

Perhaps the more interesting speculation is the reverse. If switching facilities are built for parallel computers, with low delay and high speed, why are these not more generally useful for communications switching? Perhaps we should expect that these sorts of switches, which even today are based on special VLSI, will form the basis of the local area networks of tomorrow.

#### 5.1.4. Artificial Intelligence

To this point, artificial intelligence (AI) has not been deeply influenced by the existence of communications. This field, like many others, has benefited from the move to distributed interconnected work stations, but the basic algorithms of AI are in principle independent of the computing base that executes them. There are, however, at least three interesting speculations about AI communications.

The first speculation is that an AI processor, like the brain, may be partitioned into specialized parts for reasoning, input processing, and so on. These parts have to be connected in real time if a unified facility is to be built. Current networks do not seem an obvious match with these sorts of requirements, but a future network with highly parallel high-speed flows might be much better suited.

A second and more interesting speculation concerns communication among AI processors, and between AI processors and humans. A human in isolation is much less effective than a human among others. This ought to be true of AI processors as well as humans. Humans have communication channels of varying speed: the low-speed written word, the medium speed of speech, and the very high speed of vision. Of these, vision is probably the most potent. Will this also be true of AI processors?

A critical limitation of the human brain is that it has no output device to match the eye. One of the most difficult communications for a human to do is to convey to another human a visual image. In general, it requires drawing the image, and those who spend their time with images, such as architects, spend much of their time in this very slow output process.

There is no reason why an AI processor needs to have this same limitation. If it seems appropriate, high-speed image generation could be an integral part of all AI processing. Transfer of these images might well be the basis of communication from AI processor to human or among AI processors. Perhaps most important, humans may come to use an AI image generator to facilitate human-to-human communication. The computer has already revolutionized human communication through electronic mail. AI-mediated human communication may be the next revolution, and images will probably be at the center of this revolution.

The final speculation about AI and communication relates to the representation of knowledge within the computer. We must assume that at some point in the progress of AI an attempt will be made to codify, in computer-accessible form, the bulk of human knowledge. It seems clear that the gathering as well as the use of this information will be a distributed process. What will be the communications requirements for this effort, especially if the knowledge base must be accessed in real time as part of some ongoing computation?

### 5.1.5. Teleconferencing

Electronic mail has clearly shown that the computer can greatly facilitate communication between humans. Real-time communication is an area that the computer has yet to enhance. There are many roles the computer can play in real-time communications or teleconferencing. It can coordinate the conference, automate the visual aids, and archive the conference. The communication requirements for a teleconference are very exacting if moving images of the participants are to be transmitted in real time. Perhaps the most important role of computerized communications is the coordination of these images together with the speech and the visual aids of the conference. This is an application area now being actively explored.

### 5.1.6. Image Processing

Both AI and teleconferencing have generated requirements for high-speed communication because they require images as part of the data. Over the next decade, images ought to become a basic part of all user interaction. To understand the potential impact of images within the computer, consider a comparison with text processing today. Twenty years ago, computers were tools for only writing and running programs. Out of the need to edit programs came more general tools for manipulating text. These tools have transformed the computer into the word processor, which is probably the single most important application in making the computer generally useful. General tools for manipulating text are now an expected part of almost all user interfaces and applications.

Today, we see image processing as text was seen 20 years ago. Image processing belongs in specialized domains, such as cartography, medical research, printing, and space exploration. Even though images are perhaps the most accessible form of communication to humans, there is no use of them in general interfaces. There are two reasons for this. First, the processing of images requires high performance in the display, the processor, and the communications system. The cost of this has been insurmountable in the past, but does not seem as serious an issue now. Second, systems programmers have not been interested in image systems, because images are not a part of programming, and systems programmers most often build systems for themselves. If we choose, we could build a new generation of systems, in which images were a central and ordinary part of the environment, with image filing, image display, and image processing done with the general tools of the system. Such a system would facilitate a new generation of applications, and would probably broaden the accessibility of computers to the same degree that word processing did.

The relevance of communications to image processing is obvious. The bandwidth requirements for image systems, based on specialized systems in existence today, are from one to two orders of magnitude greater than for systems that handle text. A typical image is 1 to 10 Mbit, so retrieval in one second requires as much as 10 Mbit/s. Advanced applications today have requirements that exceed this speed, and even 100-Mbit networks do not seem sufficient for some existing applications.

### 5.1.7. Digital Document Distribution

If images become one of the normal modes of computer data, then the computer and communication substrate can serve to distribute much of the material now sent as printed matter. Magazines and advertisements, for example, both require good quality images. Advertising represents a potential source of revenue to bring image distribution systems into existence. Once they exist, there is a possibility of online publishing as a viable alternative to the printed material of today.



## 5.2. FUTURE NETWORK ARCHITECTURES

Each of the technological applications we have discussed has the potential to require speed in excess of what the architectures of today can deliver. Many of these applications require delivery with low latency as well as high speed; and most of them would benefit from and thrive in the context of a network with wide accessibility, so that there could be a broad base of participation in the application.

What, then, will be the characteristics of this next generation network? The following is a speculative consideration of what that network might be.

### 5.2.1. Network Size Estimates

If this new network is to connect computers of the sort found today in offices and homes, then a reasonable size estimate might be a few million attachment points on a nationwide network. If, however, the network is to provide the same sort of ubiquitous access as the phone system now does, then the network must match the phone system in size; that is, it must have a few hundred million attachment points.

However, it would be most unwise to postulate a future computer network architecture based on some particular estimated maximum size. Rather, systems of this size must be based on the ability to scale in size, so that the architecture can grow almost without bound. It is critical that this new architecture not be trapped by a restrictive size assumption, so that it is not outgrown before it is fully deployed.

### 5.2.2. Network Speed Estimates

There are two aspects to network speed: the speed of the underlying components and the delivered speed to the users of the network. In the consideration of technology above, we postulated media speeds of 1 to 100 Gbit/s. More interesting, perhaps, is some consideration of the speeds that specific applications will require.

In the section on image processing, application speeds of 10 to 100 Mbit were postulated, depending on the degree of compression of the image. The same is true of digital video, where the uncompressed speed is about 100 Mbit. Typical disk speeds today are 15-25 Mbit/s, so if we extrapolate this out 10 years, remote disks might effectively use 100 Mbit or so. If we allow room for unexpected growth, it seems reasonable to postulate a maximum application speed of a few hundred Mbit per second.

A more interesting speculation concerns the minimum speed that this network will deliver. This consideration may not be an obvious one, but it is critical. There are two sorts of networks that might be imagined. One network would have only one sort of attachment point, with the same range of services available everywhere. The other sort of network would have a range of attachment points, with different services available depending on the particular version of the attachment point selected. If every end point of the network were to deliver exactly the same sorts of service, then application design would be greatly simplified, because every application would work equally well at each attachment point. However, in a network in which every end point can deliver 100 Mbit to the application, it is difficult to imagine how this network could extend, for example, to mobile hosts such as cars or briefcases. Cellular radio does not seem to scale to 100 Mbit per attachment point.

If the network would permit a variety of attachment services, then end points such as mobile hosts could more obviously be integrated into the system. However, the advantage of this increased

ubiquity could be minimal. Although the attachment points could be widespread, a given application, for example image transmission, could be capable of operation at only one sort of attachment point. This would result in a network that would be no more widespread than the first sort of network would have been.

Taken to an extreme, this sort of network ends up with a number of sorts of attachment points, each tailored to a specific application, and the network itself becomes a number of application-specific overlays sharing a common trunking capability. This is a common form of networking today. For example, the trunking capacity of the telephone system is used to transport both voice and video, but the attachment methods for television and telephone seem totally different.

To the extent possible, it should be a goal of this new architecture to support attachment methods that are independent of application, and to permit all of the services of the network to be accessible from all attachment points. It is almost certainly true that this goal must be compromised to some extent. Nonetheless, it should be the goal.

### 5.2.3. Resource Management

Compared to networks today, this new network must be much more sophisticated in the management of its resources. This is true for two reasons. First, it must allocate its resources among applications with very different requirements. One application may demand reliability, another controlled variance of delivery. These various sorts of data flows must be multiplexed together in such a way that the needs of each are met. Second, these requirements must be adjusted in the case where demand exceeds supply. For example, a file transfer can be slowed if there is excess demand, while a flow for real-time speech can only be throttled by terminating it altogether or compressing it. No network of today has resource management tools sophisticated enough to combine together these diverse sorts of requirements.

### 5.2.4. Addressing Dynamics

Most networks today do not cope well with attached devices that move around from place to place. The phone system, for example, is conceptually a network for connecting people, but the addressing scheme addresses phones, and if a person moves from phone to phone, the system provides no support in redirecting the call. The phone number describes a device that is not assumed to move. Mobile phones are only a partial and very special case exception to this consideration.

In contrast, the next generation network should provide addressing and routing to permit attached devices to move from one attachment point to another. If a computer is moved, or is attached to multiple attachment points for reliability, the network should understand and cope. If a person moves from one location to another, connections to that person should be properly routed.

The need for this requirement may not be obvious. It derives from the fact that communication will more and more be between computers than between humans. Humans are very effective at coping with misdirected connections. Often conversations occur in which further phone numbers are passed across the network. The form of these conversations is not constrained in any way; humans function well in the absence of such constraints. Computers do not. If this sort of address management is to be done in a network of computers, it must be architected as part of the system.

### 5.2.5. Organization Structure

It is very uncommon for a system of the size proposed here to be constructed and operated by one organization. Networks of today, such as the telephone system or the television distribution system, are composed of regions managed by separate entities. These entities must cooperate (by conformance to standards) at the same time that they compete. This multi-entity structure is a tremendous source of complexity within the system, because standards must be developed for any aspect of the system that spans the regions. Forcing functions to be realized through standardized interfaces makes system evolution very cumbersome and risks compromise of system function in the process of setting standards.

It is necessary to assume that the future network here postulated will similarly be composed of many regions run by distinct entities. Businesses are likely to operate private networks, in a manner similar to how private phone systems are operated today. Trunking capacity will be provided by competing carriers, unless the current attitudes on regulation are reversed. A management structure to cover community distribution would have to be developed.

It is possible that the multiorganizational nature of this network will be the single most difficult issue in its design. Since the network will offer a broad range of services, compared to networks of today, a broad range of standards will have to be developed. However, standards can only be developed when the technical aspects of the problem are understood, and many of the problems envisioned for this network are far from being solved.

### 5.2.6. Data Routing

An excellent example of a difficult problem that gets more difficult in a multiorganizational network is the routing of data within the network. In the network of the future, routes will have to be selected based on a number of considerations: the type of service needed, current network loading, the costs of various routes, and the security of those routes. Each of these matters can be quite complex, but to combine them is much harder still, because the issues do not all arise at the same level of the network architecture. Traffic loads, for example, are visible at a very low level in the system, within the switching nodes themselves, while security policy is a very high level matter, depending on what application and agent is using the connection.

The network described here must contain mechanisms to permit these diverse considerations to be folded together into a coherent routing architecture. Routing will no longer be a low-level matter for the switching nodes, but will be negotiated among all levels of the system. At a minimum, this suggests that routing decisions will be complex and expensive to make, so they should be computed only when necessary. Simple datagram networks of today compute routing decisions on each packet; this will certainly not be reasonable in the future.

### 5.2.7. Network Security

Most networks today do not include effective security mechanisms. Rather, they depend on the devices at the end points to enforce security. The phone system of today, for example, connects any two phones together; it is up to the people using the phones to determine if the call should be permitted. This level of security will not be adequate in the future; it will be necessary for the network and its clients to share in the enforcement of security policies.

The network must enforce security because the computers attached cannot be expected to do so. More and more, computers are becoming dedicated to single applications or persons. There is no reason why this sort of computer should be engineered to be strongly resistant to malicious attack, and the cost of adding these features to the computer are substantial. Attaching such a computer to a network will put it at intolerable risk unless the network itself can offer some protection. In short and as a practical matter, if this network is to succeed it must take on the burden of security controls.

### 5.3. PROJECTS FOR TODAY

If there is a requirement for a new generation of communications network, then we must identify the projects that we should undertake today to bring it into existence. Several projects seem relevant and are examined here.

#### 5.3.1. New Switching Facilities

This document postulates that neither packet switching nor circuit switching alone will suffice for the next generation of computer networking. In addition to advances in older switching technologies, new switching architecture candidates should be proposed and explored.

One possible approach would be to bind resource commitments to individual data flow, to ensure that bandwidth and switching capacity are available to support the quality of service required by the flow. This prebinding of resources is somewhat similar to circuit switching and should permit the same level of performance, while at the same time permit the dynamic multiplexing of packet switching. In other words, there may exist a hybrid of circuit and packet switching with the good features of both.

Another approach would depend on the fact that the nature of the network management required in a local area network is substantially less than that required in the network as a whole. If all hosts are attached to the network via a local access net, then the host need not concern itself with many of the complexities of the network architecture; the switching node connecting the access net can perform many of these on behalf of the host. If the switch can act as agent for the host, then the problems of interfacing can be greatly reduced.

These proposals must be realized in a way that permits other important goals to be maintained: robustness, security, dynamics, and manageability. Ideally, the actual switching modes of the system would be hidden from the application, which would result in the specifications of the type of service required. This would permit the network to select the most appropriate way of achieving this service: packet switching, circuit switching, or something altogether different.

#### 5.3.2. Large Networks

The issues of size and organization raised above require that basic algorithms for addressing and routing be rethought.

We could attack these problems with the techniques used in the phone network, where addresses reflect the hierarchical nature of the network itself. However, this approach does not permit the flexibility in addressing that we believe is needed to support multicast, mobile hosts, partitioned networks, and nodes with multiple attachment points. Instead, a scheme based on logical addressing and source routing might prove more fruitful. A preliminary design suggests that, given the memory costs of today, such a scheme would be practical.

### 5.3.3. Network Management

A network of this size will require management tools very different from the ones we know how to build today. The tools of the phone system are not suitable; the difference in the quality of service between the phone system and this network means that new sorts of tools will be needed. The problems that must be solved range from fault isolation to resource allocation. The latter, in particular, must be addressed as a central part of the design of the switching architecture.

### 5.3.4. Multiorganizational Networks

Even in data networks today, there are many functions that are not understood well enough to standardize. Examples of these functions are fault isolation, rate optimization, or routing among mutually suspicious regions. The network of the future will have many more of these problems. If we are to have a practical multiorganizational network, these must all be understood, and the solutions standardized. The complexity of some of these problems suggests that new approaches to standards may be required, with much more dynamic negotiation occurring between regions. It is important to study now in an organized way how networks can be assembled out of separately managed regions.

## 5.4. SUMMARY

In this paper we have proposed bringing into existence a new generation of computer network. This new network will provide performance for the applications of the next decade, and will provide widespread accessibility in the manner of the current phone system.

There are a number of applications that might benefit from this network: supercomputers, parallel processors, AI research, teleconferencing, and image processing. It is not clear which of these will have the most compelling need, but taken together, they constitute a strong argument for the network. We must remember our experience with networking today, wherein electronic mail arose as a new and unexpected application, and we must recognize that we may not be able to predict the most important outcome of a new network effort. If we are in the position to build a facility that extends by two or more orders of magnitude the capabilities of today, we ought to do so, if only to provide fertile ground for new applications to take root.

Given that we should build such a network, is it practical? We believe that it is. Fiber optics and VLSI provide the new building blocks for this next generation network. The fiber can provide the speed at a low cost. The VLSI can harness the speed and make it accessible to the application. Together, these two technologies make the network possible.

There is a final requirement if such a project is to succeed: there must be people with ideas who want to build it. In fact, proposals already exist for the high-speed architecture and the large size network, and researchers are excited and challenged by the possibility of building such a network. What is needed is a consensus that the effort is needed, and a commitment to see the work done.

We believe that it is critical that this research be undertaken now. The urgency arises from the need to explore some of the technical issues well before the time when standards must be set. The open negotiation of standards is a slow and inflexible process; the research results must be available in a timely manner if they are to contribute to the success of the standards-setting process. Even a lead time of a decade for this research may not be sufficient.



## 6. ADVANCED SYSTEM SOFTWARE FOR SUPERCOMPUTERS

---

Dieter Fuss and Richard W. Watson  
Lawrence Livermore National Laboratory\*

### Abstract

*Advanced systems software development for supercomputers falls primarily into two areas--distributed computing and multiprocessing (parallel processing). Requirements for advanced systems software development for distributed computing and multiprocessing can be viewed at four levels: the operating system, the application run-time environment, tools or utilities, and command language interfaces. A review of the requirements placed on systems software development and future software needs for distributed computing and multiprocessing at each level are discussed in this paper.*

### 6.1. INTRODUCTION

A distributed computer system is a collection of processor memory-pairs connected by a communications network and logically integrated by a distributed operating system. The communications network may be a widely geographically dispersed collection of communication processors or a local area network, or a combination of both. A special case of distributed computer systems is a multiprocessor system. A multiprocessor system may be a collection of processor memory-pairs loosely coupled via a very high speed communications network (bus), or a collection of processors tightly coupled via shared memory, or a combination of both.

Distributed and multiprocessor computer systems have been the subject of much research. Many prototype systems have been, or are being developed at university, commercial, and government research institutions. It is impossible to review here all of the projects and systems. Instead, in this paper we will discuss the motivation for the forces that are driving development of advanced system software for supercomputers and detail the requirements placed on that development if the goal is to maximize the efficiency of man and machine.

### 6.2. ADVANCED SYSTEMS SOFTWARE REQUIREMENTS

There are two major forces driving software systems development for supercomputers:

- The demand for computer networks and distributed computing needed to increase user productivity, maximize efficient use of resources, and enhance resource sharing.
- The demand for multiprocessor hardware architectures needed to increase simulation performance of physical phenomena and other large-scale calculations.

We will first review the motivations for developing distributed computing and multiprocessing and define our usage of the term "systems software." Then, we will outline some of the requirements placed on systems software development by distributed computing and multiprocessing.

---

\*Work performed under the auspices of the U.S. Department of Energy by Lawrence Livermore Laboratory under Contract W-7405-Eng-48.

### 6.3. DISTRIBUTED COMPUTING

In the next decade, supercomputers, communication networks, mass storage systems, input/output (I/O) systems, and intelligent terminals will all be viewed as elements of a distributed computer system having resources managed by a single distributed operating system. The price/performance revolution in microelectronics, the development of cost-effective and efficient communications networks, and the development of powerful, highly interactive workstations have set the stage for this integrated view of large-scale computing.

Supercomputers no longer operate most economically as standalone self-contained entities. Two motives drive the need for distributing function--user requirements and system requirements.

#### 6.3.1. User Requirements

Scientists, engineers, and administrative personnel are a precious resource; their productivity must be increased. The architecture of large applications is becoming more modular; and applications, or parts of applications, could be distributed to improve user and system economics, utilization, performance, interactivity, and reliability (for example, separation of pre- and postprocessors from the computation-intensive applications or separation of interactive and computation-intensive functions in editors). Moreover, workstations have brought the power of mainframe computers to the desktop along with high-performance, high-resolution color graphic displays, mouse-driven input, and highly interactive software. By using workstations and network connections to other services one can provide

- improved human-machine interface techniques (windows, mouse, menus);
- improved interactivity;
- improved access to (and sharing of) tools, information, and databases; and
- improved communication between people through electronic mail and easier information sharing.

#### 6.3.2. Systems Requirements

Large computer centers are in a constant state of evolution. Most large supercomputer centers have several mainframes attached to a network that is also interconnected to mass storage system, printers, terminal connections, etc. One of the major system-software issues is how best to create a system-integration framework that can live and evolve over decades and that allows integration of this heterogeneous hardware and software. Unfortunately, the software that supports the sharing and access to such heterogeneous processing, storage, I/O, and other special services was not designed as a whole. Rather, each subsystem was designed independently and then, over time, each subsystem was connected to the larger system and programmed to understand every other subsystem--i.e., the software was typically developed as a point-to-point system.

Although it all works, functionally it may have limitations, and it is very difficult to modify or add capabilities in such an environment--each change impacts every other subsystem. What is needed is a single network-wide view of the total hardware/software environment with a standard interface to any subsystem--i.e., a software network bus. Such a software system would minimize the cost and impact on users and system software when replacing or adding new services or systems by

- improving modularity and distribution to support extensibility and evolution;
- allowing different parts of a system to evolve at different rates appropriate to changing user requirements and to technology change; and



- improving maintainability, reliability, and availability through distribution of functions.

## 6.4. MULTIPROCESSING

The second major force driving developments in systems software arises because of the constant need for more processing cycles balanced with larger memories and high-performance I/O. The history of performance gains in supercomputers is well known. As technology shifted from vacuum tubes to transistors and then to integrated circuits, basic machine cycle times went from milliseconds to nanoseconds--an improvement of six orders of magnitude--in less than 40 years. But such tremendous improvements are now a thing of the past; the speed of light is becoming the barrier to faster cycle times.

Newer technologies, e.g., gallium arsenide (GaAs) devices, are likely to shrink the distances signals will have to travel; but even with this technology, machine cycle times are not likely to become less than a nanosecond or two in the next 5 years, which would represent about a factor of 10 speed-up over 10 years.

Although the speed-up in cycle times resulting from technological advances is the clear common denominator in improving supercomputer performance over the years, changes in architecture have also played an important role. Prefetching instructions and operands (lookaheads), providing registers to "hide" slow memory, pipelining or duplicating functional units to achieve more results in fewer cycles, and vector instructions are all examples of architectural changes that have contributed to the overall performance improvements of supercomputers.

Over the last few years, parallel processing has become the architectural alternative to the lack of faster cycle times. Utilizing the concept that replicating hardware is easier and cheaper than developing new technology, vendors have turned to duplicating vector units and even whole processors to allow computation of a single problem to proceed in parallel. The NEC Information Systems, Inc., SX-2 computer allows up to 16 parallel vector operations to be performed simultaneously. The Cray Research, Inc., X-MP/48 and the CRAY-2 each have four complete CPUs attached to a common memory to allow either four independent jobs to compute in parallel or to allow one or more jobs to "multiprocess," i.e., to compute pieces (tasks) of the job in parallel. Both Cray Research and ETA Systems, Inc., are developing eight-processor systems. Floating Point Systems, Inc., recently announced a system that is expandable to 16,384 processors; such a system is called "massively parallel."

The use of multiple processors to increase **throughput**, i.e., to improve overall system throughput by overlapping the execution of multiple unrelated tasks, is rather straightforward. The use of multiple processors to decrease single job **turnaround**, i.e., concurrently to execute subparts of a single logical task, is much more complicated. Parallel processing to achieve a reduction in the real run time of a single job, whether it is vector processing on multiple vector units or multiprocessing on multiple CPUs, requires a greater understanding of both the application and the computer architecture on which it is to execute.

Technological advances in the capacity of memory chips have led to the availability of memories on supercomputers that are two orders of magnitude larger than 5 years ago. Memory sizes are at 1 to 2 Gbyte today and will reach 8 Gbyte in the next 5 years. This is fortunate in the sense that supercomputer calculations have frequently been limited by lack of memory in the past, but, coupled with the fact that single processor speeds are not increasing at nearly the same rate as the memory sizes, it implies that very large codes must be multiprocessed to achieve reasonable turnaround.

Moreover, codes using all of the memory will either have to use all the processors or all but one processor will sit idle.

Multiprocessor systems are also of interest for operating systems in the supercomputer environment. For example, the dedicated systems supplying mass storage and other shared services can potentially utilize multiprocessor systems to improve reliability and provide incremental increases in processing power as their loads grow.

## 6.5. SYSTEM SOFTWARE

Distributed computing and multiprocessing have many common issues. Whether the potential they represent can be realized will depend heavily on developments in system software.

What is system software? System software is that set of programs that integrates and allocates (fairly and securely) the hardware and simplifies the application writer's and end user's view of the system by providing interfaces (languages) and logical resources more appropriate to application needs than those provided by the raw hardware (for example, files instead of raw disk drives). It is convenient to view the system software environment as consisting of four levels.

### Level 1: The operating system

The operating system provides two main services.

- First, it turns a collection of interconnected heterogeneous hardware/firmware/software resources into a set of abstract objects or resources (e.g., processes for computing, files for storage, directories for human-oriented object naming, clocks for timing, accounts for cost recovery, and messages for communication). This service should support object intercommunication, naming, sharing, protection, synchronization, and error recovery in as uniform a manner as possible.
- Second, it multiplexes and allocates the resources mentioned immediately above securely and fairly among concurrent cooperating, competing, or possibly mutually suspicious computations.

### Level 2: Application run-time environment

This is a set of libraries or programming language extensions providing access to the operating system services and supporting other services.

### Level 3: Tools or utilities

These are editors, compilers, debuggers, file utilities loaders, library managers, etc.

### Level 4: Command language interfaces

These allow access to the system services through workstations, terminals, or batch interfaces.

## 6.6. REQUIREMENTS FOR ADVANCED SYSTEMS SOFTWARE

We will now discuss some requirements that the two dominant trends in supercomputer architecture and environment imply for system software in each of the four layers outlined above.

Most existing operational or about-to-be-released commercial software systems for supercomputers are derived from developments of the late 1960s and early 1970s, when supercomputers were single-processor standalone systems. These system software environments were not originally designed with communication, distribution, or multiprocessing in mind. Also, most were not designed with adequate

mechanisms for privacy and security. When computer networking began to evolve in the 1970s, supercomputers were integrated into these environments in a very loosely coupled way, basically supporting remote terminal access and file transfer via explicit user command. The communications protocols were usually incompatible across applications. No official or de facto communication protocols standards existed. Further, the operating systems were primarily single-processor batch systems, and they only recently evolved to support single-processor interactive use.

## 6.7. DISTRIBUTED COMPUTING REQUIREMENTS

The main requirement of the entire spectrum of system software for distributed computing is to provide support for programs, programmers, and terminal users that gives a **single network-transparent uniform view** (hiding heterogeneity) of **all resources in the system**. The resources include such things as processes, files, and databases. Network transparency means that programs, programmers, and users should not have to be explicitly aware of whether a resource being used is local, remote, or made up of distributed parts. Neither should they program differently or use different terminal procedures depending on resource location. Of course, performance or economics may depend on location; therefore, it must be possible explicitly to control the location of a resource or program module when desired.

A second requirement of the systems software architecture is extensibility; because the life of the total environment is long, spanning decades, and the technology and the user needs are constantly changing.

A third set of requirements is a result of various networks and computers being administered by different organizations. Each organization may want to control who has access, what programs and services can be accessed or run, the time during which their physical resources are available, the quantity of resources available, and the accounting for access or resources. Groups of systems may also cooperate to provide a distributed file, process, or other service, thus requiring distributed resource management, access control, and accounting policies.

There may be mutual suspicion between systems. This can result because of differences in administrations, physical security, or local-component operating systems. This implies that the basic protection mechanisms must not be built assuming correct or secure operation of all systems and networks. Mutual suspicion implies modularity of protection or access control, often called the principle of least privilege access, whereby a program has access to just those resources needed for its function and no others. This principle is also important in limiting error propagation. It should be easy to pass access rights dynamically between programs. Multilevel secure communications and computing must also be supported at some installations. This implies the requirement that all objects in the system (e.g., packets, files, processes, communication links, etc.) be labeled with a security level and that an appropriate security policy be implemented assuring storage, transmission, and access to only appropriately labeled objects.

### **Level 1: The operating system for distributed computing**

The future operating system architecture required to support the above environment must create the illusion that there is a single system supporting network-transparent access to uniformly defined resources. The term used to describe such an operating system is a network or distributed operating system spanning all the nodes. The total operating system will consist of all the local component operating systems at each node working together according to agreed communication and interface standards.

**Level 2: Application run-time environment for distributed computing**

It must become as straightforward to program and debug a distributed application as a centralized one. The main requirements that some combination of language or library features must support are the following.

There must be the ability to invoke local and remote functions and pass parameters in a way that is independent of the caller's and callee's locations. In particular, user programs should not have to deal with the underlying system heterogeneity of the communicating systems or the forms and encodings of network messages when invoking remote functions.

At the user's level, the programs should access local and remote services and programs with the same mechanisms.

It must also be possible to debug distributed and concurrent multiprocessing applications as easily as local or single-process applications.

**Level 3: Utility support for distributed computing**

The main requirement is that utilities should be able to access both local and remote resources. For example, a text editor on a workstation should be able transparently to access and edit named files anywhere on the network, or a debugger should be able to control distributed modules and process structures. These distributed modules may be running concurrently and be programmed in different languages. Features are also needed to examine message contents, trace messages, and so forth.

**Level 4: User interface support for distributed computing**

One important way to increase user productivity is to improve the responsiveness of, and ease of, communication between man and computer. Event handling, program swapping, and context switching are expensive on supercomputers. Therefore, an important direction of software development is toward moving the user-interface software into scientific workstations that support high-resolution display screens, multiple windows for concurrent human activity, and the development of highly interactive interfaces. User requests will be specified and be parsed on the workstation using interfaces defined around graphical representations and pointing devices; remote calls will then be made to application functions on the back-end supercomputers. This will require more modularity in application development and carefully separating the user-interface function from the computation and data management aspects of the application. It will also require developing tools in the workstation to support distributed user-interface specification and implementation. In addition, remote invocation (remote procedure call) protocols and windowing software must be supported.

**6.8. MULTIPROCESSING REQUIREMENTS**

Users need increasingly detailed and accurate simulations of physical phenomena that must be run to completion in reasonable wall-clock time. Because the industry is reaching limits on switching device speed and signal propagation delay, performance improvements in the future must come through parallel architectures. These architectures will range from the fastest possible multiple processors with large high-bandwidth shared memories to separate processor/memory modules coupled by high-speed backplanes and to relatively autonomous distributed processors with specialization of function. In addition to the applications that require the fastest possible sets of individual processors, there are applications that could meet their performance and economic needs by multiprocessor systems made up of tens to hundreds of low-cost, mass-produced, and slower VLSI processors and memory.

### **Level 1: Operating system for multiprocessing**

There are two basic hardware multiprocessor communication models, i.e., multiple CPUs connected to a shared memory and separate CPU/memory modules communicating over a high-speed bus via messages. Many different architecture variants with features of both will abound. These variants may include several levels of nonshared and shared memory. An important operating system design area is to support the communication between, and manage the imbalances in speed of, these memories. The choice of the hardware model will affect many of the base operating system requirements. For example, in the shared-memory model, the same operating system code, data structures, and tasks may be shared among CPUs; while in the separate memory/processor module model, independent operating systems run on each module. Use of shared memory implies an efficient mechanism for synchronizing shared memory access; while in the message passing system, special firmware and operating system support is required for very efficient message passing.

Both the user application and the operating system on a multiprocessor node should be capable of concurrent execution (multithreaded) using all available processors. For reliability and maximum performance, there should be no master/slave operation with the operating system running in only one processor.

There must be inexpensive mechanisms to spawn and switch control among concurrent entities. Ideally, costs for these should be on the order of those for a procedure call. Mechanisms are also needed for concurrent entities to share memory, communication ports, and other resources, and to be able to synchronize access to these shared resources. It is desirable that processor scheduling algorithms allow concurrent entities in a single job simultaneous control of available processors to maximize memory utilization or improve message passing performance. An important design issue is how much of the multiprocessing scheduling mechanism to place in the operating system and how much in libraries and in user codes.

Even if one only wants an architecture to support the distributed environment outlined earlier, many of these same characteristics are needed in the design of service programs supporting concurrent service access from many locations.

### **Level 2: Application run-time environment for multiprocessing**

The main language or library need is to support algorithms that contain concurrency and synchronization. This can be achieved by developing new language features or using existing languages with library support. In addition, there is the need to develop language processor mechanisms (compilers) to detect parallelism within programs not explicitly written for multiprocessors and automatically to produce code that can make effective use of multiple processors.

Libraries written for single-processor systems cannot usually be safely used in shared-memory multitasking applications written for multiprocessor systems, because they do not protect their shared data structures to assure mutual exclusion. Further, the single-threaded libraries themselves do not utilize multitasking where performance might benefit.

Debugging multiprocessing codes is complicated. The order of execution of parallel tasks is unpredictable and, therefore, it is particularly difficult in a timesharing or a multiprogramming environment to get the repeatability useful in debugging. To achieve this repeatability requires support within multitasking libraries or language run-time systems to force repeatability--for example, control execution from traces. This is an important area needing research and development.

Another issue of importance is determining the level of granularity with which to perform multitasking. If done in small granules such as loops, the term multitasking is used. If done at larger algorithmic levels, the term multiprocessing is used. The particular hardware architecture used can dictate the size of granularity for which multiprocessing is economic. With microtasking, a few simple directives supported by an appropriate library can be used to convert a program originally written for a single processor quickly to one that can achieve significant performance gains. Use of multitasking requires much more work on the part of algorithm and program developers, but it may yield better performance or be useful with a wider range of hardware architectures. The pros and cons of these two approaches are under active study.

Many existing applications make use of signaling or interrupt facilities to stop current program execution and force control to specified routines for exception handling. The semantics of such mechanisms are unclear in a multiprocessing environment; therefore, new language, library, and system mechanisms are needed to support or replace these existing facilities.

### **Level 3: Utility support for multiprocessing**

The main new types of utilities needed to support multiprocessing are enhanced interactive debuggers and utilities to aid or automate creation of multitasking programs. The enhanced debuggers should control the sequence of multiple task execution, analyze the source and order of synchronization events, and examine individual task state. The utilities to aid or automate creation of programs are needed to help analyze codes for areas where performance gains may be possible and safe for multiprocessing, show where shared variable conflicts may exist, measure and display degrees of overlap (that is, the number of processors actually in use in parallel at a given point), trace the order of task execution, synchronize events, and so forth. This is an area requiring extensive new development.

### **Level 4: User interface and command line interpreter design**

Terminal users of many single-processor systems can stop the execution of their programs, inquire of their status, restart them, and perform other related control actions. The meaning of status inquiry and these other actions is currently unclear in a multiprocessing environment.

It is also unclear how status on the multiple tasks may be obtained by the system when, as is common, the scheduling of the tasks and of their state is being handled in libraries or in run-time environments within a single operating system supported process; and, thus, the scheduling is invisible to the operating system. Related to the remarks made in the section above is the question of what types of information or displays may be most useful in showing the activity of these multiprocessing jobs.

## **6.9. CONCLUSION**

### **6.9.1. Distributed Computing**

In the next few years we will see increased growth of local and wide area computer networks linking supercomputers to a range of smaller general and special-purpose systems, including workstations. Initially, these networks will mainly be used for remote terminal/workstation access to other systems and file transfer. Network-based resource sharing will be common. Distributed computing involving application structures with two or more cooperating processes running on separate processors will become a reality. International communication protocol standards will begin to appear in commercial products, although they will probably not satisfy some of the efficiency and security requirements for distributed computing (see Section 6.10).

The spread of distributed computing will be largely dependent on development of appropriate industry standards that will yield a distributed operating system. Standards are needed to support remote function invocation and standard operating system services (file, directory, process authentication, accounting, etc.). Such standards are being worked on, but the process is inherently slow.

De facto standards will probably appear earlier than official standards. For example, the Sun Microsystems Network File System protocol is becoming a de facto industry standard protocol for distributed file access. It will be used in a variety of implementations to facilitate a tighter file system coupling between supercomputers and other computer systems.

There is active research in distributed computing within universities, some computer vendors' research departments, and government laboratories. The AT&T UNIX system is becoming something of a de facto standard. Unfortunately, many of the UNIX architecture features do not distribute easily, UNIX has serious security problems and, while multiprocessor versions of UNIX exist, they place substantial limitations on application structures (see Section 6.11). These limitations can only be overcome by changes in its architecture and major redesign of its implementations. In general, commercially available systems will continue to be weak in their distributed system capabilities and be incapable of supporting multilevel security, as their security capabilities will be weak in general. Only in leading-edge laboratories capable of system software development and integration are limited distributed computing applications and supporting development tools likely to appear in the near future.

### 6.9.2. Multiprocessing

Although more research in algorithmic and software issues of multiprocessors still needs to be done, all system software for supercomputers will eventually be modified to support some form of multiprocessing. The differences between systems will be in the extent to which the operating systems, libraries, and utilities themselves can be multiprocessed; their efficiency in batch and time-shared operation; the ease with which they can be extended to support larger numbers of processors; and their features to support multiprocessed application development.

From an applications viewpoint, it is not cost effective to move a code to a new machine architecture unless the move can be done fairly automatically or a sizable speed increase is realizable which will, in turn, allow new computational problems to be solved. Vendors of traditional supercomputers are counting on research that will make code portability to a multiprocessor fairly automatic and therefore cost effective. Vendors of massively parallel supercomputers count on their prediction of very large performance gains to encourage programmers to go through the effort of rewriting codes in a language designed to support multiprocessing.

## 6.10. INTERPROCESS COMMUNICATION ARCHITECTURES FOR DISTRIBUTED OPERATING SYSTEMS

A distributed operating system is likely to be built on the client/server model.<sup>1-4</sup> In this model, abstract resources of a particular type--such as processes, files, directories, and timers--are managed by processes called servers. Processes called clients access the resources by sending request messages for service to the appropriate server and normally expect replies containing data or service confirmation in turn. A given process can act in both client and server roles at different times. For example, a directory or process server may be implemented as clients of a file server.

There may be many servers for a given type of resource; for example, files distributed within the environment may be managed by different servers. A given application is likely to dynamically access resources residing on many systems while running. These interactions are transaction oriented; that is, a client process issues a request and the server issues a reply. No additional conversation need ever take place. We believe the communication protocol architecture should support this style of transaction-oriented interaction with minimal overhead.

In general, requests and replies require communication between heterogeneous systems. Therefore, issues arise about how to encode standard data types used in function- and parameter-encoding between the heterogeneous systems. Different systems, for example, may use different encodings of data types such as integers, floating point numbers, character strings, etc. Even systems of the same type need ways to serialize structured data types for network communication. To assist in translation between systems, a network standard is required for parameter and data encoding.

Unfortunately, existing protocol standards such as those of the Department of Defense (DoD) or the International Standards Organization (ISO) do not meet all the above needs. As a result, researchers in distributed operating systems have been developing their own network protocols.<sup>1-2,4-5</sup> We can briefly outline the deficiencies of the DoD and ISO protocols. The basic problem with both architectures is that they have, at the transport and higher levels, a high overhead for setting up and tearing down connections to complete a transaction as outlined earlier. Both the ISO transport and session protocols<sup>6</sup> and DoD transmission control protocol (TCP)<sup>7</sup> require three-way handshakes to set up and tear down connections. This means that in many common implementations, as many as nine packets may have to be exchanged to send one request and receive one reply. Alternatively, at some point of initialization all processes that may converse may have to have "prewired" connections, with the resulting overhead of having to know, a priori, both all the required communicating parties and about maintenance of connection states. The implementations are also often large and perform slowly.

The higher level file transfer protocols of both protocol suites (DoD and ISO) do not support low-cost file access. In the case of the DoD suite, the file transport protocol (FTP) can only be used for whole-file transfer. While the ISO file transfer, access, and management (FTAM) protocol supports partial file access, it requires very high overhead connection establishment and is, thus, not likely to be widely used in distributed operating systems. Further, the DoD protocol suite does not define standard data encodings for request/reply parameters and data. The ISO presentation level does define such standards.

The DoD and ISO protocol suites do define a connectionless or datagram network-level protocol on which transaction style transport and higher level protocols meeting the needs of distributed operating systems can be built. Thus, while the services provided by the transport and higher levels of the DoD and ISO protocol suites will be useful in limited, loosely coupled terminal access, whole file transfer, and electronic mail applications, they do not currently support the standards necessary for efficient distributed operating system implementation.

## **6.11. PROBLEMS WITH UNIX AS A DISTRIBUTED OPERATING SYSTEM FOR SUPERCOMPUTERS**

The UNIX operating system<sup>8</sup> was designed over 20 years ago to be used by a small number of people for interactive program development on a minicomputer; it was not designed for supercomputers or for distributed computing or multiprocessing. Security/privacy was not a concern, reliability of the hardware was assumed not to be a problem, and high-performance computation was not a goal.



Retrofitting UNIX to overcome these deficiencies requires major architectural and/or implementation changes. We list here some broad areas of concern.

- The monolithic structure of the UNIX kernel does not lend itself to extensive modifications.<sup>9</sup>
- UNIX has no unified model of abstract objects and functions to operate on them. Processes constitute one world, files another, and newer objects such as message queues and semaphores yet another. This presents difficulties in developing extensions to UNIX.<sup>9</sup>
- There is no location independence when accessing resources. Knowledge of distribution must be embedded in all parts of the system that access resources. That means each part must check if the resource is local or remote and, if remote, invoke a separate kernel--kernel or library--library mechanism. Performance in the UNIX systems that have tried distribution tends to be quite poor.
- Security and access control are very primitive and limited. There is no multilevel security policy; there are no explicit security labels on objects. UNIX has many published security holes.<sup>10</sup>
- Because of blocking I/O and weak synchronization mechanisms, the UNIX kernel cannot be multitasked without significant redesign, and it is difficult to multitask within a UNIX user process.<sup>11</sup> The latter requires new asynchronous system calls and other process-related system changes.<sup>12</sup>
- There is no warm restart or job recovery facility. Accounting and batch job facilities are minimal.

## 6.12. REFERENCES

1. G. T. Almes, A. P. Black, E. D. Lazowska, and J. D. Noe, "The Eden System: A Technical Review," *IEEE Trans. in Software Engineering* SE-11, 1, pp. 43-59 (1985).
2. M. B. James and R. F. Rashid, "MACH and Matchmaker: Kernel and Language Support for Object-Oriented Distributed Systems," Central Michigan University, Mount Pleasant, Michigan, Computer Science Department report (1986).
3. R. Sandberg, et al., "Design and Implementation of the SUN Network File System," Sun Microsystems, Inc., report (1985).
4. R. W. Watson and J. G. Fletcher, "An Architecture for Support of Network Operating System Services," *Computer Networks* 4 (1980).
5. J. G. Fletcher and R. W. Watson, "Mechanisms for a Reliable Timer-Based Protocol," *Computer Networks* 2, pp. 4-5, pp. 271-290 (1978).
6. "Information Processing Systems--Open Systems Interconnection--Transport Protocol Specification," International Standard Organization report ISO/DIS 8073 (1984).
7. "DoD Standard Transmission Control Protocol," Information Sciences Institute report RFC-761, Marina del Rey, California (1980).
8. D. M. Ritchie and K. Thompson, "The UNIX Timesharing System," *Communications of the ACM* 17, 7, pp. 365-375 (1974).
9. R. F. Rashid, "Threads of a New System," *UNIX Review* 4, 8, pp. 37-49 (1986).
10. G. R. Grossman, "How Secure is 'Secure'," *UNIX Review* 4, 8, pp. 50-63 (1986).
11. M. D. Janssens, J. K. Annot, and A. J. Van De Goor, "Adapting UNIX for a Multiprocessor Environment," *Communications of the ACM* 29, 9, pp. 895-901 (1986).

12. T. W. Hoel, "CRAY-2 Unicos Kernel Enhancements," *Cray User Group Proceedings*, pp. 18-20 (Spring 1986).

## 7. LOCAL AREA NETWORK TECHNOLOGY WITH EMPHASIS ON GATEWAYS AND DIGITAL SWITCHES

---

L. W. Amiot, J. M. Boyle, P. C. Messina  
R. J. McMahon, and L. Winkler  
Argonne National Laboratory

### Abstract

*No longer are scientists or engineers locked into a single computer in order to get their work done. Local, national, and even international networks are emerging in great profusion, and making available to these users a vast and diverse set of hardware, software, and database resources. The potential benefit of sharing these types of resources must be weighed against the cost of interconnecting the user's work environment with the shared resource. The success of the resulting network is measured not only in the cost of its connectivity but also in its ability to meet the functional needs of the user as well.*

*The purpose of this paper is to discuss the technological issues of interconnecting resources via local area networks and the impact of their connection to wide area networks. The basic networking strategy presented is based on a concept of a hierarchy of local area networks where each network layer is connected to the next by a gateway or bridge. Each local area network would be based upon a hardware/software technology that best meets the functional requirements of its users. The gateways provide the means of coupling local area networks of various differing technologies or networks that are under different organizational control. While the ideal situation would be to have a single network standard that would meet all of the user's functional requirements, that is currently not the situation, and local area networks of varying technologies and varying manufacturers exist.*

### 7.1. NETWORK FUNCTIONAL REQUIREMENTS

We believe, from the user's point of view, that there are four major types of functional requirements for computer networks:

- electronic mail,
- file transfer,
- terminal access, and
- distributed computation.

We also believe that users require these services not just within their local network, but across hierarchies of local networks and wide area networks as well. In particular, access to these functions must be transparent through network gateways.

If satisfactory service is to be provided, each of these functions induces its own requirements on the design of networks and their interconnection mechanisms. In the following discussion, we classify these requirements with respect to the following network properties:

- **Latency**--the delay in passing a message across-a network
- **Speed**--the rate at which information can be transmitted
- **Bandwidth**--the total capacity of the network

### 7.1.1. Electronic Mail

Electronic mail involves the sending of a message to a user or a computer, possibly one that is several networks away from the sender's computer. From the user's point of view, the message should arrive in a timely fashion, say within a few minutes, if all required networks are operating. When a required network path is down, the message should be queued for delivery when the path becomes available. Information about the disposition of the message (for example, return receipt requested) may be desired. Unlike a telephone call, receipt of the message must take place without any action on the part of the recipient. Electronic mail messages are typically small. However, they provide a convenient mechanism (because simultaneous interaction of the sender and recipient is not required) for the exchange of manuscripts or computer programs whose size may be 100,000 bytes or more. Additionally, users require convenient access to a "Yellow Pages" server to locate electronic mail addresses.

Existing electronic mail schemes, such as those used on the ARPANET, are basically satisfactory. Difficulty of interchange and addressing across dissimilar networks are currently weak points, but emerging mail standards can be expected to solve many of these problems independently of network advances.

Electronic mail demands only limited speed and responsiveness from the network. However, to be useful it does require a store-and-forward capability somewhere in the network. For electronic mail to be useful, the networks used must be reliable in the sense of being available without long outages.

### 7.1.2. File Transfer

File transfer involves a user initiating a request to send a file or files to, or to receive a file or files from, a remote computer. From the user's point of view, he would like the files to be transferred in real time and as quickly as possible. He would like easy access to the remote computer, while at the same time he would like to be assured that his files are protected from unauthorized access by others. To be effective, file transfer requires some limited terminal-like access to the remote computer in order to find the directory and file to or from which transfer is to take place. The volume of data being transferred can be quite high, 10 Mbyte or more, and such volumes need to be transferred in at most a few minutes. Of course, the data must arrive without error and with suitable character translations having been made. At the same time, binary or other machine-specific data must not be corrupted in the transfer.

Existing file transfer mechanisms [for example, a file transfer program (FTP) on the ARPANET, and especially remote procedure call (RPC) among UNIX systems] are basically satisfactory. Higher speed and greater network bandwidth would improve usability.

File transfer demands high network speed to minimize the time a user must spend effecting the transfer. Network bandwidth must be adequate to handle the anticipated number of simultaneous file transfers. Latency is not a serious problem except when interactive listing of directories or files is taking place.

### 7.1.3. Terminal Access

Terminal access involves using a remote computer from a terminal or workstation as if it were a local resource. From the user's point of view, it should not be apparent that the computer being used is thousands of miles away. Currently, a typical "message" from a user to the remote resource in terminal access mode involves a few characters--in many cases, a single character. Messages from the remote computer to the terminal are larger on the average, although single character messages will occur when echoing in full-duplex mode. In the future, larger messages may be more common, for example, for graphical output from the remote computer.

Existing terminal access services, such as *Telnet* on the ARPANET and especially *rlogin* on UNIX systems, are basically satisfactory within local area networks. Currently, however, wide area networks do not adequately meet the requirements discussed in the next paragraph, resulting in user dissatisfaction.

Terminal access currently requires only relatively low network speed (tens of kilobits per second) and limited bandwidth. However, it sets stringent limits on the latency of the network. Many applications must run full-duplex (or generalizations thereof, such as graphical response to user input). Network latency of even several tenths of a second can seriously impair usability. Achieving low latency may preclude the use of satellite links in wide area networks and places stringent demands on latency in gateways. Similarly, sufficient bandwidth must exist in the network if low latency is to be maintained while other high-bandwidth functions are taking place. When connecting local area networks to wide area networks, it may be necessary to think in terms of two parallel networks: a low latency network for terminal access, which would use overland and underwater links, and a higher latency network for the other functions using satellite links.

### 7.1.4. Distributed Computation

Distributed computation involves running a multiprocessing user program on two or more computers on a network simultaneously. We expect such applications to be of increasing importance over the next few years. A few initial applications are already in place. One is remote access to file systems (for example, via Network File System). In this application, a file system on a remote computer can be mounted and made to appear to be a local file system on the user's computer. Remote files are used by sending individual blocks of the file over the network as required, rather than by transferring the entire file at one time. Another distributed application is remote display of graphical output, in which graphical output is sent from a large supercomputer via the network to a workstation for display. We also think that distributed editors and window systems will soon be commonly used. Within the next 1-2 years, distributed simulations and other application-oriented computations will become common. Such distributed computations will require the relatively frequent exchange of messages over the network to maintain synchronization among the distributed processes.

Existing distributed computations are either prototypes or special-purpose (for example, Network File System). Thus, the adequacy of present networks to support general-purpose distributed computation is largely unknown.

Distributed computation can be expected to place severe demands on latency, speed, and bandwidth. Remote file access requires substantial bandwidth and low latency, since the file may be used interactively. Graphical output makes similar demands, especially in regard to latency, because the user may need to interact with the program running on the supercomputer that produces the output.

Distributed simulations and other application programs will require high speed and low latency, and possibly high bandwidth, to effect synchronization among parallel processes via message passing.

## 7.2. LOCAL AREA NETWORK PROTOCOLS

It would be nice if one network protocol existed that would provide for all of the functional requirements described above and if all manufacturers agreed on such a standard for access to their equipment. Unfortunately such a standard does not exist and many manufacturers have gone their own way in the development of network protocols. The Open System Interconnect (OSI) reference model being developed by the International Standards Organization (ISO) might be just such a standard. Unfortunately, we are still a couple of years away from completion of the specification and another couple of years before it becomes widely integrated into manufacturer's equipment. For the time being, at least, we will find it necessary to deal with several such network implementations when integrating a campus-wide network. Among those of greatest importance are DECnet, Transmission Control Protocol/Internet Protocol (TCP/IP), and IBM Systems Network Architecture (SNA) and NJE.

### 7.2.1. DECnet

DECnet is a networking protocol developed by the Digital Equipment Corporation (DEC). It has been implemented principally for DEC computers: PDP11s and VAXs. It allows interconnection of DEC computers using synchronous telecommunication links, x.25 links, or Ethernet cables. DECnet utilizes their proprietary DDCMP as a data link protocol. File transfer, terminal passthrough, electronic mail, and distributed computation are all supported up to and including the application level. The protocol is full function but limited to DEC equipment and limited in its network addressing capabilities.

### 7.2.2. TCP/IP

TCP/IP was developed as part of the Defense Advanced Research Projects Agency (DARPA) internet protocol suite. It is a Department of Defense standard and is an integral part of Berkeley BSD UNIX operating systems. The protocol is primarily designed to run on Ethernet but can use HDLC as a data link protocol on telecommunications lines. The protocol is full function and supports terminal access, electronic mail, and file transfer. TCP/IP is general in nature with several standards described at the application level: FTP (file transfer), Telnet (terminal access), and SMTP (mail protocol) are examples. The protocol has been developed for many manufacturer's computers and operating systems. Its integration into applications programs, such as for distributed computing, although limited, is beginning to appear.

### 7.2.3. IBM SNA and NJE

IBM developed its Systems Network Architecture (SNA) to define communication between processes running in a host and between processes in different host computers. The intent of the architecture was mainly to encompass IBM and IBM-compatible computers. Synchronous and asynchronous telecommunication links are covered as well as channel-to-channel adapters for colocated computers. The principal communications protocol is SDLC, synchronous data link control. The architecture covers IBM access method levels down through line protocols with only application level interfaces defined. Terminal access is defined and the provision for transferring files exists. Application level file transfer, electronic mail, and distributed computation are left to application program development. IBM has defined a general-purpose interprogram synchronous protocol for distributed applications known as LU 6.2.

NJE refers to a compatible set of software developed for various IBM operating systems that provides file transfer and electronic mail. Those applications predate SNA and make use of the older IBM binary synchronous protocols. They are now supported under SNA as well.

### 7.3. INTEGRATION OF LOCAL AREA NETWORKS ON CAMPUS

Local area networks tend to be installed on an organizational basis but be limited by geographical constraints. They are usually installed within a building on a departmental, divisional, or programmatic basis. This development is brought about because of funding considerations and a desire to control the use of resources. Scarce resources should be made available to the users who represent the program that paid for them. Overutilization by others leads to slow response or poor availability. Control ensures that priority of use can be enforced. Control ensures that adequate operational and maintenance procedures can be put in place. Control of the network also allows for a greater measure of security of information to be provided. These reasons tend to argue for small local area networks each under local programmatic control.

The need exists for users within one organizational unit to communicate with and utilize the resources of other units. Granted the need for interorganizational communication is less than intraorganizational communication, but the need for terminal access, distributed computation, file transfer, and electronic mail between resources of various organizations is still important. One of the tasks of the network implementor is to provide connectivity between the resources of one organization and those of another without loss of control.

Since the trend is to develop local area networks based on organizational units, the need arises to interconnect the local area networks of one organization to those of another. The result is a campus-wide or laboratory-wide network. On the lowest level, one can consider the campus-wide network to be a single high-speed bus that serpentineing around campus and interconnecting to each local area network, while on the highest level we can consider the local area networks as each having a gateway where the gateways are interconnected by a network in a hierarchical arrangement. It should be noted that the implementation of a campus-wide bus, which directly interconnects local area networks, tends to dilute the control users have over their local networks, while the use of gateways tends to maintain control. However, as discussed below, the use of gateways has its drawbacks as well.

### 7.4. LOCAL AREA NETWORK COMMUNICATION BACKBONES

In choosing a local area network, three basic issues must be addressed: the network's transmission media, its topology, and its access method.

#### 7.4.1. Transmission Media

The choice of network architectures depends to a great extent on the amount of bandwidth required. Four types of transmission media are currently implemented: twisted-pair wire for low speeds (up to 1 Mbit/s), baseband coaxial cable for medium speeds (up to 10 Mbit/s), broadband coaxial cable for high speeds (up to 500 MHz aggregate total), and fiber optic cable for very high speeds (up to 1 Gbit/s).

Baseband technology uses a single high-speed digital channel that is used primarily for data communications. Baseband is implemented in a bus topology over which signals travel in both directions. Broadband is a wide bandwidth local area network technology that carries analog signals at up to 500 MHz using frequency division multiplexing to divide the total bandwidth into multiple 6-MHz channels. Each channel is capable of carrying a data, voice, or video signal. Multiple data

subchannels at speeds up to 5 Mbit/s can be established to allow separate networks to share the same physical medium, resulting in a high aggregate data rate. Broadband networks are laid out in a tree topology over which signals travel in only one direction. The maximum distance limitation is 12,464 feet. Hybrid network configurations are being designed to take advantage of both technologies. Small local baseband subnetworks are established taking advantage of its high speed and easier maintenance. A broadband backbone network is used to carry information between subnetworks over greater distances, thus minimizing the number of direct connections to the broadband backbone.

The principal local area network in use today is baseband Ethernet, which can utilize coaxial cable or fiber optic cable. It is possible to construct one large Ethernet network that would interconnect all of the campus resources without the use of bridges or gateways, but the use of Ethernet does have distance limitations. Ethernet comes in two types: thin and thick Ethernet cable. Thin cable has the most severe distance limitations (5,260 feet with repeaters), while a network of thick cable could be 9,194 feet with repeaters. Thick and thin can be intermixed with thin cable connecting to the thick by a simple passive connector. Electronic devices, such as the Siecor FiberLAN products, allow a number of Ethernet fiber optic segments to be interconnected in a physical star arrangement in order to cover larger geographical areas than a serpentine cable can cover. The maximum cable distance between any two devices that wish to communicate would still be limited to 9,194 feet, however. The Star Wiring Center eliminates the need for repeaters. Other devices are available that allow two remote Ethernet segments to be bridged over large distances by the use of fiber optic cable and drivers. Ethernet runs at a 10-Mbit/s rate, which is shared by the devices connected to the cable.

#### 7.4.2. Network Topology

A network's topology is the physical and logical arrangement of its nodes. The three basic local area network topologies are bus, ring, and star. In a bus network, the nodes are arranged along a single length of cable. A tree is a complex bus in which the cable branches. All broadband and many baseband networks use a bus or tree topology. Today, the bus and tree technologies are the most prevalent due largely to the Ethernet community, which has established that particular bus technology as a standard.

In a ring topology, the nodes are arranged along the transmission path so that a signal passes through one node at a time before returning to its originating node; the nodes are arranged to form a closed circle. IBM's Token Ring Network is an example of a ring network. A star network has a central node that connects to each station by a point-to-point link. AT&T's Starlan network is an example of a star network. All transmissions in bus and ring networks are broadcast. All signals transmitted on the network pass all the network nodes. Each node looks at the signals, recognizes its address, and copies only its own signals. In star networks, signals are switched through the central node to the proper receiving node.

#### 7.5. ACCESS METHODS

A network's access method is the technique used to distribute the right to transmit among participating nodes. The right to transmit is only an issue in broadcast technologies, where a number of nodes share a single data channel on which all nodes receive and on which any node can transmit. The access method is the network's means of controlling traffic. A network's access method is an important factor in determining its performance. In a Carrier Sense Multiple Access (CSMA) network performance degrades as the likelihood of collision increases. A CSMA network operates more efficiently when nodes transmit long individual messages than when nodes transmit many short messages. Because the Ethernet protocol is contention based (CSMA), average throughput rates could



be reduced by a factor of 10 or more when large numbers of devices are contending for its use. Token-passing networks perform better under uniform heavy loads than CSMA networks. The number of nodes is an important factor affecting performance in token-passing networks. Other networks are now available that run at significantly faster rates. The Proteon, Inc., implementation of the token passing access technology allows 80-Mbit/s network throughput to be realized. The University of Illinois has implemented a ProNET-80 high-speed campus backbone network interconnecting smaller Ethernet-based local networks.

Local area networking standards, developed by the Institute of Electrical and Electronic Engineers (IEEE), are in place for the physical and data link layers of the OSI reference model. These layers define the physical and electrical characteristics of the transmission medium and the medium access management. Standards include the IEEE 802.3 standard for baseband CSMA (Ethernet-type) networks; the IEEE 802.4 standard for token passing bus networks; and the IEEE 802.5 standard for token-passing ring networks.

New network technologies are being developed that will push the network data rates even higher. These are based on both technological (fiber optics) and architecture (token passing ring networks) advances. One such development, the Fiber Distribution Data Interface (FDDI), promises data rates of 100 Mbit/s. The counter rotating ring technology utilizes distributed token access on a fiber optic based transmission medium. Work on this standard is being done by the Accredited Standards committee X3T9 Technical Committee for I/O Interfacing and the X3T9.5 Task Group for LANs Over 50 Mbit/s.

## 7.6. GATEWAY TECHNOLOGY

Gateways provide the means of connecting networks in order to deliver packets of information from one network to another. In their simplest form, they become bridges that connect two networks of identical design and relay all information that appears on one network to the other. As the level of intelligence of the bridge increases, it performs functions such as filtering. Only information from the originating network that is addressed to the destination network is relayed through the bridge. Unnecessary internetwork traffic is thus minimized.

Full-fledged gateways provide another layer of intelligence as compared to the bridge. They not only bridge two or more networks but they provide network routing as well. As information is received from one network by a gateway, it examines it and determines for which of multiple networks the information is destined, and then routes it to that network over the appropriate interface. The routing is based on internal tables maintained within the gateway and on information contained within the received packet.

Vitalink provides the TransLAN product to bridge multiple Ethernet or IEEE 802.3 local area networks together so they to appear as one single network using high-speed DDS or T1 connections. The Vitalink TransLAN product provides continuous adaptive learning of the network for easy installation and dynamic reconfiguration. TransLAN simultaneously supports local area networks running TCP/IP, Xerox Network System (XNS), and/or DECnet. The TransLAN product has selective filtering of data, only allowing data addressed to another network to pass through the TransLAN, providing efficient local area network utilization and network security. The Vitalink TransLAN hardware works especially well in broadcast satellite networks. A sophisticated management system is provided for monitoring traffic and supporting configuration services.

The Proteon ProNET-Linkway is a gateway designed to allow two computers operating the same network protocol to communicate even if they are located on networks that utilize different architectures or transmission media. Interface hardware is available for connection to local or wide area networks including the ProNET-4, ProNET-10, ProNET-80, Ethernet, ARPANET, and synchronous communications lines. The ProNET-Linkway software currently performs packet forwarding for IP and XNS. In the future, Proteon plans to add software to its gateways to implement DECnet, IBM's SNA, ISO, and MAP packet forwarding, and hardware interfaces to IEEE 802.5 and T1.

Some gateways have the ability to update their routing tables dynamically by exchanging information with another gateway. The DARPA internet model defines an autonomous system as a collection of networks of identical design interconnected with gateways that communicate with each other using a well-defined Interior Gateway Protocol (IGP). An example of an IGP is the Gateway-to-Gateway Protocol (GGP) used by the internet network. A protocol, the Exterior Gateway Protocol (EGP), has been defined to allow gateways between autonomous systems to exchange routing information with each other.

### 7.6.1. NJE Gateways

Using the above gateway definitions, we can examine the various protocols discussed earlier. NJE essentially utilizes a gateway at each node. Each is a store and forward node that examines its internal static routing tables to determine the destination network to which to forward the stored information. It should be noticed that software has been developed for some computers, such as the VAX NJE emulator developed at Argonne National Laboratory, that makes them end nodes and not gateways on a NJE network because they do not have store and forward capabilities. The NJE gateway code is simple and reliable (it stores information until it can be delivered), but the accuracy of the routing is only as good as the effort that goes into manual maintenance of each routing table in the network.

### 7.6.2. DECnet Gateways

In a DECnet network, routing nodes route packets from one node to another. For very large networks, DECnet implements area routing. In a single area, all routers can route packets to the intra-area nodes. In multiple networks, only certain routers can route packets between inter-area nodes on the network. DECnet calls the intra-area routers level 1 routers and the inter-area routers level 2 routers. With the latest DECnet Phase IV software, a 16-bit number is used to address each node; the most significant 6 bits define the area number and the least significant 10 bits define the node number within an area. Thus, up to 63 areas each with 1,023 nodes can be addressed.

### 7.6.3. Internet Gateways

Given a set of systems, all of which implement the internet protocols, it is rather easy to implement a campus-wide hierarchy of local area networks. Hosts can be connected on a departmental level using Ethernet as the communications medium. An array of software and hardware is available to implement the internet protocols: Sun and Ridge workstations, DEC VAX computers running Berkeley UNIX 4.2 or 4.3, VMS running Wollongong 2.3, or IBM VM or MVS operating systems with Spartacus K200 hardware and KNET software are all examples. Internet software is also becoming available for personal computers, principally IBM PCs at this time: Wollongong WIN/PC, Spartacus KNET TCP/PC, Network Research Corporation Fusion-PC, MIT PC/IP, and Excelan EXOS 8051 are all examples.

The technology currently exists to integrate the departmental networks into a campus-wide network. As previously described, campus-wide Ethernets or Proteon networks can be installed to be internet compatible. Departmental networks can then be bridged or connected via gateways. DEC VAXs running Berkeley UNIX 4.2 or 4.3, VMS running Wollongong 2.3, or Sun workstations are examples of systems providing gateway capabilities.

#### 7.6.4. Protocol Conversion Gateways

Of course, not all networks follow the internet protocols. Digital Equipment Corporation's DECnet and IBM's SNA and NJE are examples. While each of these can easily be integrated into campus-wide networks using standard vendor products, integrating them together and with the internet networks, so that campus-wide communications among dissimilar networks is possible, is a difficult task. Gateways must be constructed to interconnect them, which not only route packets but which reformat packets as well. Conversion of information from one protocol to another can involve conversion of control information (handling of "return receipt request" for mail protocols is an example) as well as code conversion (ASCII to EBCDIC is an example). We can see that if we had N different dissimilar networks, we would have full communication between all networks. If we could define a standard network protocol, we could then construct gateways that convert from the source protocol to the standard protocol and then to the target protocol. Each of the dissimilar network protocols would have to have a conversion to and from the standard. This approach however only requires N different conversions to be coded. Extra overhead would, however, be required to do the double conversion to and from the standard. Much work is still to be done to optimize gateway designs of this type and to spell out the tradeoffs between coding and performance.

Spartacus, Inc., offers a K200 front-end processor and KNET/VM networking software to allow IBM mainframes to communicate over Ethernet with IBM and non-IBM workstations and computers. The K200 interfaces the Ethernet to the IBM channel and the KNET software interfaces TCP/IP protocols to IBM protocols (Telnet to IBM 3270, for example).

Digital Equipment's DECnet/SNA gateway is a communications processor with associated software that functions as a DECnet end node on the DECnet network and as an IBM Physical Unit Type 2 node on the SNA network. DEC software (access routines) runs on individual DECnet hosts to provide end use functions (IBM 3270 terminal emulation, remote job entry, and program-to-program communications). The gateway provides the protocol translation required to communicate between the DECnet and SNA environments.

Interlink's IBM MVS/DECnet and IBM VM/DECnet allow IBM's System/370 architecture computers to participate as peer-level nodes on a DECnet network. The two gateway products allow IBM machines to perform high-level DECnet functions such as remote file access and remote record access, submission of remote print and batch jobs, and task-to-task communication. The gateways give IBM users similar capabilities through utility programs and subroutine libraries. Bidirectional network terminal emulation is a feature of the IBM MVS/DECnet gateway product. The IBM and DEC computers are linked via an Interlink 3711 Network controller that connects to an IBM channel and to an Ethernet cable.

FlexLINK is proprietary connectivity software that allows incompatible computers (VAXs running VMS and IBM computers running VM) to communicate. FlexLINK utilities allow a user to retrieve files, start tasks, or simulate a terminal as if the user were connected to the remote computer. FlexLINK's hardware connection is made through standard vendor interfaces; the IBM Device Attachment Control Unit to the IBM channel and a direct memory access module to the VAX.

IBM's LU 6.2 protocol is likely to become the standard for non-IBM to IBM SNA communications. Gateways from local area networks to IBM SNA networks will use this protocol to achieve program-to-program applications such as electronic mail and file transfer.

## 7.7. NETWORK PERFORMANCE

Two of the most important network parameters of interest to a user are the network's latency and throughput characteristics. Terminals are getting faster as are access links to local resources. Users are accustomed to fast response and demand it over network connections. The use of file servers and shared high-quality printers is becoming more prevalent, and thus network throughput for file transfer must be high. The use of a gateway between networks, however, tends to add delay in transferring information between networks. The need to reduce overall delay argues against the excessive use of gateways in the connection of local area networks. More work needs to be done in the area of analyzing the impact on network performance across gateways and on developing gateway designs that maximize performance. The delay of a gateway, say a DECnet router, which connects nodes via dedicated telecommunication links that run at 9600 bit/s, is insignificant to the transmission time of information. As network rates increase to 10-Mbit/s Ethernet speeds, the delay through the gateway becomes appreciable. With communication backbones that are 10 or 100 times faster, the gateway technology may become limiting.

## 7.8. IMPACT OF THE DIGITAL SWITCH

Digital switches that are presently under development offer the promise of providing a campus-wide backbone network. One such switch that is currently available, the InteCom S/80 PBX, provides not only simultaneous voice/data service, but also emulates a local area network. Internally, the switch operates as a proprietary, packet switch network with distributed nodes and a centrally located switch. The distributed nodes are connected to the switch via fiber optic cable at distances of up to 25,000 feet. The distributed nodes are connected to station equipment via two-pair, twisted building wire at distances of up to 2,000 feet.

The InteCom system has stations that can be electronic voice only, simultaneous voice/data, or data only. The data can be asynchronous at speeds to 19.2 kbit/s, synchronous at speeds to 56 kbit/s, or various local area network standard interfaces. One such standard is the Ethernet standard which describes the interface between an Ethernet device and a transceiver connected to a local area network. This InteCom station allows devices that would normally connect to a transceiver on an Ethernet cable to connect to a piece of InteCom station equipment instead. The switch thus emulates an industry standard Ethernet network and allows devices connected to it to communicate as if they were connected to a real Ethernet cable.

Data from a connected Ethernet device can be transmitted at 10 Mbit/s to the station equipment where it is buffered for transmission into the switch as a packet of information. The twisted pair allows the data to be sent to the distributed node as a 980-kbit/s information stream. The information then contends with up to 15 other Ethernet devices for a 980-kbit/s, full-duplex channel on the fiber optic link, which connects to the central packet switch. There are 422 such full-duplex channels going into the switch, giving rise to a total data capacity of 1 Gbit/s.

The conversion of the protocol from an industry standard to the proprietary protocol at the station equipment makes it an easy job to develop different standard interfaces to the switch. In addition to the Ethernet device interface, an interface is available that allows a real Ethernet cable to be

gatewayed to the switch. When two such Ethernet cables are connected to the switch, the switch effectively bridges the two networks. Devices connected to one cable can communicate with those on the other as if they were all connected to the same cable. The Ethernet gateway, as implemented by the switch, performs an important function of isolating traffic on one network from the other. Only broadcast messages or messages destined for a host on the other network traverses the gateway. The gateway learns the Ethernet addresses of the various hosts on the Ethernet cable and builds a table within the station equipment to tell which packet to relay to the other network and which to ignore.

Unfortunately, an Ethernet gateway cannot always keep up with the traffic that devices on the Ethernet cable would like to transmit to it. This comes about because of two reasons: first, the Ethernet cable is capable of 10 Mbit/s while some of the paths in the switch are shared at 980 kbit/s, and second, the Ethernet cable to which the gateway is sending information may be too busy to receive it. Thus the originating gateway has two choices: it can discard packets it cannot handle and rely on the higher level network protocols to retransmit them, or it can put up a contention signal on the Ethernet cable and cause transmission on that cable to cease. The options of discard or contention are programmable in the InteCom equipment.

We can now summarize the implications of a digital switch on interconnecting local area networks if we can consider the InteCom model as being typical of the capabilities we can expect to be developed by that market segment. The digital switch can provide a campus-wide Ethernet compatible network that has a greatly larger geographical area (54,000-foot-diameter circle) than can be covered by a conventional Ethernet cable. It can provide a standard jack in every office within that area to which an Ethernet device can be connected (on the order of 8,000 devices per S/80 switch). The digital switch provides throughput capabilities that are comparable to that of Ethernet, but because of architectural differences between the digital switch and contention based Ethernet cable, one would find that either could out perform the other depending on the loading and traffic distribution assumptions. The digital switch provides the ability to isolate traffic from individual local area networks as opposed to a single large Ethernet, but the switch has the disadvantage of having to occasionally discard packets or alternatively throttle down traffic. The choice of the digital switch versus an Ethernet cable as a campus-wide network then becomes a tradeoff in capabilities that the authors believe favor the digital switch. As cable-based Ethernets are replaced with communication backbones that are 10 or 100 times faster than conventional networks, the choice will be more difficult and tend towards the higher speed backbones unless digital switches also track this higher speed technology, perhaps with digital switches designed around high-speed token-passing rings.

## 7.9. NETWORK ADDRESSING

Network addresses provide the information by which packets from an originating host find their way to the correct destination host. On a single local area network, addressing can be rather straightforward and the address space need only be large enough to cover the maximum number of hosts that will ever be connected; routing is not an issue. However, as local area networks become interconnected by gateways into a hierarchy of networks, a single address field is no longer sufficient and a minimum of a network number (area number) and host number is required.

### 7.9.1. DECnet Addressing

We are already seeing network limitations in the addressing schemes of some of the existing network protocols. The most notable limitation is that of DECnet. DECnet has an area number, host number scheme in which the area number space is 63 and the host number space is 1,023.

Because DECnet is widely used both nationally and internationally, the requirement that DECnet gateways be set up is increasing. The area and node numbering convention is causing serious management problems in dealing with wide area usage. DECnet growth is causing duplicate area number problems when interconnecting large DECnet networks. Since DECnet doesn't allow duplicate area numbers in its connected networks, network managers are having to be very careful in designing their networks and making sure they have absolute control over assigning area numbers and allowing access to the network. It is not known whether the next release of the DECnet software, Phase V, will alleviate this addressing limitation.

### 7.9.2. Internet Addressing

TCP/IP requires each host on an internet network to have a unique address. The internet protocol uses a 32-bit address field that is divided into two parts: a network part and a host part. There are three class of internet addresses: Class A allows for 128 networks each with 16,777,214 hosts, Class B allows for 16,384 networks each with 65,534 hosts, and Class C allows for 2,097,152 networks each with 254 hosts. This addressing scheme provides flexibility in assigning address to networks and allows for a large number of small to moderate sized networks. The internet protocol address specification defines a small number of networks with a large number of hosts, a moderate number of networks with a moderate number of hosts, and a large number of networks with a small number of hosts. Many organizations have found this two-level addressing hierarchy inadequate. In many campus environments it is necessary to use more than one local area network cable to cover a campus. An organization that must use more than one local area network has three choices for assigning internet addresses:

- (1) Acquire a network number for each cable. This will cause an explosion in the size of internet routing tables. Gateways are required to transport packets between hosts on different networks.
- (2) Use a single network number for the entire organization and assign host numbers without regard to the local area network a host is on. This approach requires that bridges discover which local area network a host is on by using a broadcast algorithm. As the number of local area networks grows, the cost of broadcasting grows as well.
- (3) Use a single network number and partition the host address space by assigning subnet numbers to each local area network. This approach requires some modification of the original IP protocol that not all vendors have implemented uniformly to date.

### 7.10. NAME SERVERS

As we have seen above, addressing schemes can be quite complex. As hosts are added and deleted and networks change, keeping gateway and host routing tables current can be a difficult job. While dynamic updating of gateway routing tables helps alleviate these problems, the fact that networks can be quite distant from one another and under quite different organizational control makes the tracking of changes a significant task.

The ARPA internet is a large network that is likely to grow much larger. The need to map host names to internet addresses is beginning to stress the existing mechanisms. Currently hosts in the ARPA and MILNET internet are registered with the Network Information Center and listed in a global table. The size of the table and frequency of updates are near the limits of manageability. The ARPANET is in the process of implementing a distributed database that performs the name to address mapping function and avoids the problems associated with the centralized database. The ARPANET solution divides the entire internet into a structured name space referred to as a domain name space. The domain name space is a tree structure with each branch representing a set of network objects

(hosts, mailboxes, processes, etc.) controlled by an organization. A domain name is assigned to each branch of the tree. The complete domain name specification of a network object is the path from the root of the tree to the object. By convention, domain names are read from left (most specific level) to right (least specific level) with domain names separated by dots. Name server programs are being implemented to manage and control the distributed database. A name server has complete information about and authority over a set of one or more domains. Each domain must provide redundant name servers to provide the name to address resolution service for all hosts in the domain outside as well as inside the domain.

The transition from a flat name space to a structured space with distributed data management in the ARPANET has not been fully implemented to date; however, sites must conform to the domain-style naming conventions by March 31, 1987. Thus, it is premature to speculate whether this approach is more or less manageable and what affects it may have on network performance.

## 7.11. NETWORK SECURITY

Security concerns exist in the context of networks in several ways. Networks can provide unauthorized people access to a wide variety of computing resources. After all, networks are designed to facilitate access. The need for security is clear for systems that contain classified or sensitive information. But even for systems that don't have such uses, there is the danger of malicious destruction of other people's files, stealing time on computers, stealing commercial or licensed software, and invasion of privacy.

At present, networks have little protection from such problems. The following quote summarizes the situation succinctly: "The safest policy in using networks is to assume that any network can be broken, that any transmission can be recorded, and that most can be forged."<sup>1</sup> A few new facilities have been identified that might improve the security of networks.<sup>2</sup>

Provision for transparent or user-invoked encryption should be added. This facility would need to somehow interface across different local area networks and wide area networks. New techniques may need to be developed to carry out such transmissions. Gateway systems may well serve a crucial role here.

Additional instrumentation and tools for control of networks would help prevent, detect, track down, and terminate break-in attempts. For each session or connection attempt, networks must provide information on the specific origin and path of the sessions. Tools should be available to network managers to disconnect a node that is suspected to be the origin of security threats. Session termination should disconnect the entire virtual circuit. At present, some modes of disconnect leave incoming ports enabled and vulnerable to break-ins for a period. In addition, facilities should be provided to timeout idle circuits.

While these facilities are needed on all networks, including local area networks, they also have to be implemented in such a way that they function across networks as well. One security problem that is specific to Ethernet-based local area networks is that a node can masquerade as another, thus permitting unauthorized access. New technology must be developed to prevent this.

## 7.12. INTERACTIONS BETWEEN CAMPUS AND WIDE AREA NETWORKS

The implementation of the next step in the hierarchy of networks requires the interconnection of various campus-wide networks on a national and international scale. The communication backbone that provides this connectivity is the wide area network. The wide area network must supply the same functionality as does the campus, local area network: namely, electronic mail, file transfer, terminal access, and distributed computation. We judge that communication among collaborators is currently the most important function required of intercampus communication and thus the importance of electronic mail and file transfer.

In the near future, it will be necessary to connect all of the network protocols that have been implemented on campus to the wide area network. The choice then is to either build several wide area networks each of an architecture particular to the architectures of the local area networks it interconnects (a DECnet wide area network to interconnect DECnet-based networks and an internet wide area network to interconnect TCP/IP-based networks as examples) or to have a campus gateway capable of connecting to a heterogeneous set of local area networks on one hand and a standard wide area network on the other hand.

The first alternative is simple to implement using the same network protocols that were used to implement the local area networks. The cross-country links can be composed of telecommunication links that range from 9.6 to 50 kbit/s on common carrier circuits or faster using microwave. The use of 1.544-Mbit/s T1 links is starting to emerge, and nationwide fiber optic links at even higher speeds can be expected in the future. The problem with this alternative lies with the duplication of cross-country links required to implement the different wide area network protocols.

The second alternative is more economical in the use of communication links but poses the problem that a gateway must be constructed at each campus that is robust enough to handle the conversion of all of the campus networks to the wide area network standard and yet still be able to repacket information quickly and efficiently. Off-the-shelf hardware and software is not generally available to handle these conversions. It is thus important that a standard for wide area networks be agreed upon and that gateways be developed that economically and efficiently provide the necessary code conversions required of the local area/wide area networks gateway.

The most economical solution of all would be for all networks, both local and wide area, to utilize the same network protocols. Since it is generally perceived that the industry is heading towards the OSI architecture as a standard, it should be expected that this standard will be utilized for wide area networks as well. The gateway between the local area and wide area networks will then become a simple routing gateway and not require protocol conversion. Research in the area of developing OSI local and wide area networks and gateway implementations is required as the standard is being finalized.

Another concept for a national network that is currently under study is the Integrated Services Digital Network, ISDN. The ISDN is a set of broad technical recommendations for a common user interface to a digital telecommunications network that will be designed to carry voice, data, telex and videotex, and broadcast audio/video. It will provide local digital loops of at least 144 kbit/s with high-speed digital channels being available at approximately T1-carrier speeds (1.54 Mbit/s). The ISDN standard is under study by the Consultative Committee for International Telephony and Telegraphy (CCITT). At its current stage of development, the ISDN specification is little more than a map for future technical development. Although many aspects of the proposed ISDN have been implemented, no working model for the entire specification exists. ISDN can only be thought of as a



telecommunications solution for the twenty-first century. However, because of the international importance of the CCITT, this standard is likely to have far-reaching effects on data communications. The extent to which it will be compatible with the ISO standard is not known, but the importance of the ISDN and its compatibility with OSI warrants future involvement and study by government researchers.

### 7.13. TOPICS FOR FUTURE STUDY

This paper has outlined some of the more important networking protocols and networking issues that the network designer must deal with today. An examination of the issues provides insight into areas where further network research and development would be appropriate. These areas include

- monitoring tools at every level of the network hierarchy to provide statistics necessary to identify network bottlenecks,
- analysis of the impact of gateways on performance and work on gateway designs that maximize performance,
- work on gateways that provide protocol conversion between network protocols,
- development of local and wide area implementations that use the OSI standard,
- evaluation of networking alternatives by prototyping commercial products,
- development of high-speed communication backbones for both local and wide area networks,
- work on local and wide area network standards committees
- studies of the efficiency of name servers in managing network addressing,
- development of encryption standards and techniques,
- development of network management tools to prevent, identify, and track network security violations, and
- studies of the economic and technical impact of the ISDN on future networking activities.

### 7.14. CONCLUSIONS

This paper has discussed the state of the art of local area networks that can presently be found on campuses and at government laboratories. Alternatives for implementing local area networks, gateways between local area networks, connections to wide area networks, and the impact of digital switch technology on local area networks were presented. Areas for research and development were highlighted as topics for future study. It is expected that in addition to research by universities and government agencies, much of the future local area network development will be done commercially. It thus becomes important for us to take part in committees that are developing networking standards that will guide industry and in prototyping network hardware and software to further evaluate alternatives for future use.

### 7.15. REFERENCES

1. J. S. Quarterman and J. C. Hoskins, "Notable Computer Networks, Legal Issues," *Communication of the ACM*, **29**, 10 (October 1986).
2. F. M. Atchley and P. C. Messina, "Summary of Recommendations of Network Security Discussion Group," *First Annual Workshop on Energy Research Computing*, Lawrence Berkeley Laboratory report LBL-22460, CONF-8610138 (October 1986).



## 8. THE ROLE OF GRAPHICS WORKSTATIONS IN SUPERCOMPUTING

---

**T. A. Lasinski**  
NASA Ames Research Center

### Abstract

*The use of graphics workstations in the National Aeronautics and Space Administration's Numerical Aerodynamic Simulation Program is discussed. The NAS workstation environment consists of over 30 Silicon Graphics Inc. graphics workstations networked to a CRAY-2 and other computers via both HYPERchannel and Ethernet. Present application software is described with particular emphasis on distributed graphics between these workstations and the CRAY-2. A video tape illustrating a typical distributed application is available. Present and desired data rates are discussed. An attempt is made to extrapolate this experience over the next few years.*

### 8.1. OVERVIEW OF THE NAS SYSTEM

The National Aeronautics and Space Administration (NASA) established the Numerical Aerodynamics Simulation (NAS) Program at its Ames Research Center with three principal goals in mind:

- to provide national computational capability to insure continuing leadership in computational fluid dynamics (CFD);
- to act as an agency pathfinder in the integration and use of large-scale computer systems through the use of state-of-the-art computer hardware and software technologies; and
- to provide a strong research tool for NASA's Office of Aeronautics and Space Technology.

The NAS system is described in greater detail in Reference 1. Here only those features of NAS relevant to graphics workstations and supercomputing are discussed. Figure 1 shows all the computers associated with the NAS system and how they are networked together. The CRAY-2 has a 4-ns clock and 256 million 64-bit words (2000 million bytes!) of main memory. The four CPUs can obtain a peak speed approaching 1000 MFLOPS for vector calculations. Connected to the CRAY-2 are many Silicon Graphics "IRIS" workstations. Both HYPERchannel and Ethernet are used. An ever increasing number of remote sites are connected to NAS over terrestrial lines ranging from 56 to 224 kbit/s. For the last 2 years, NASA's Langley (LaRC) and Lewis (LeRC) Research Centers have had IRIS workstations connected to the NAS system.

### 8.2. WORKSTATION REQUIREMENTS FOR THE NAS PROGRAM

The NAS graphics workstations were chosen through competitive procurement in the spring of 1984. The principal requirements were as follows:

- stand-alone CPU with VAX (11/780) performance
- real-time, dynamic graphics capability
- large disk space
- good, direct communications to the Cray computers

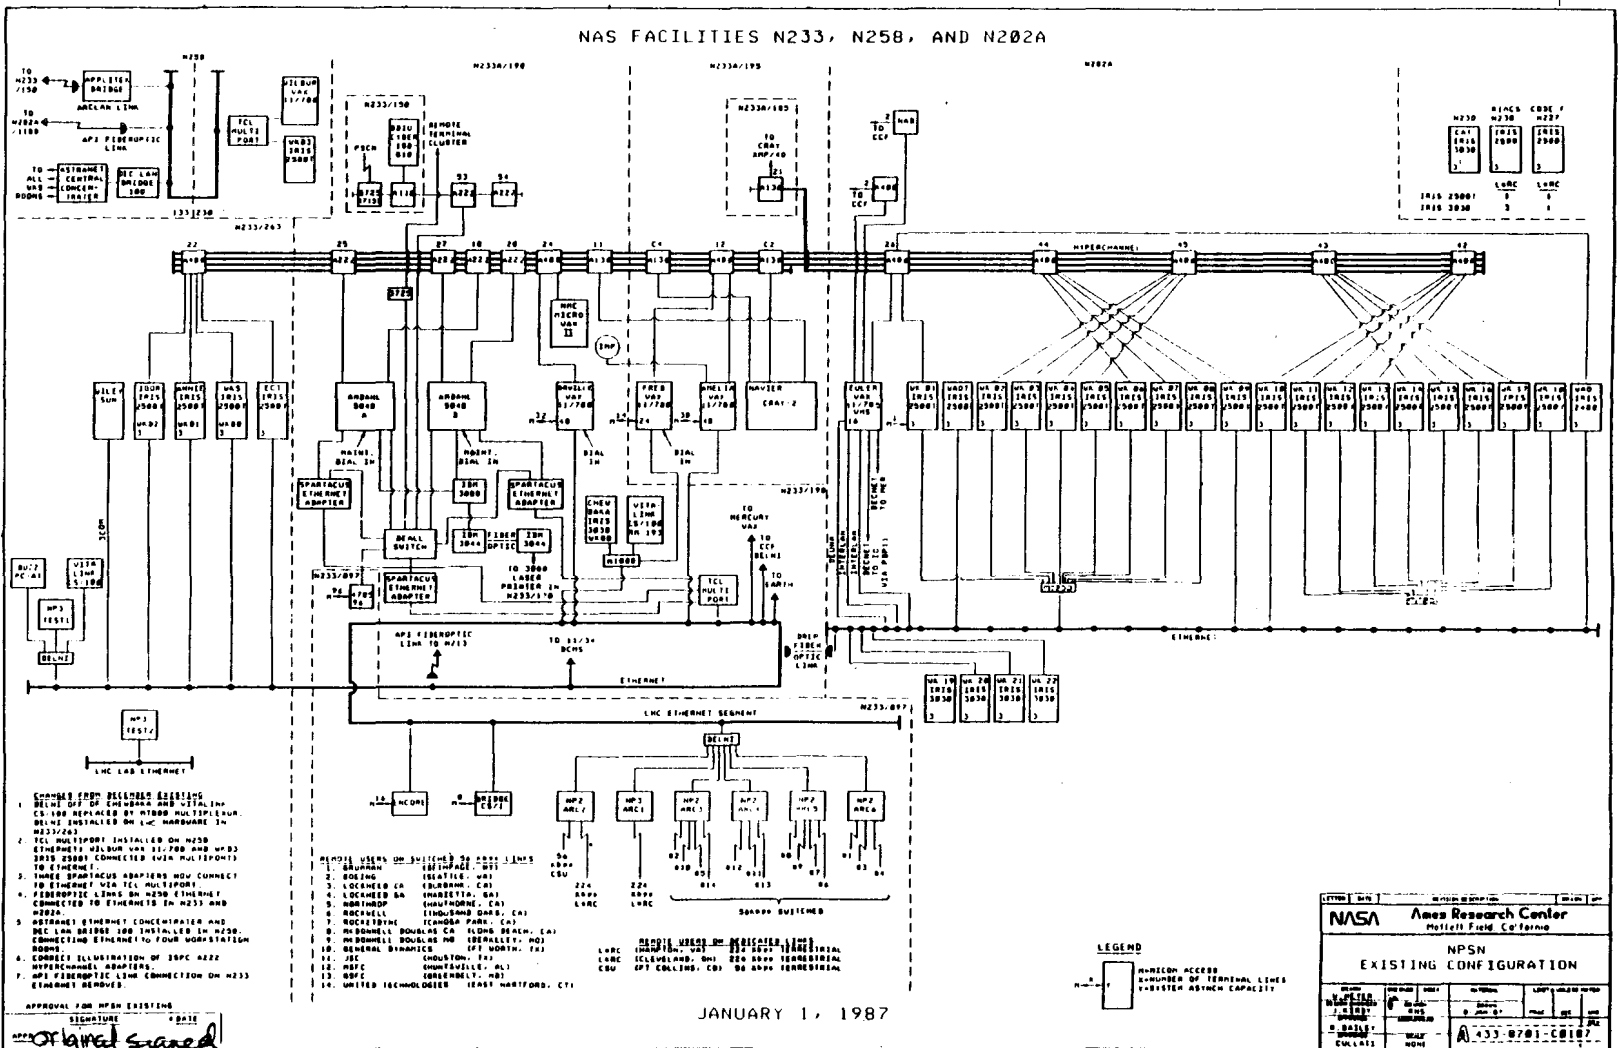


Figure 1. NAS computer system  
FCCSET Report to Congress

- capability to make video animations
- tools to make "journal quality" output

These requirements are no doubt common to all potential workstation users. The second one was of particular interest in that it distinguished this workstation from the vast number of commercial workstations available then. While bit-map graphics were commonplace among workstations, the requirement called for the capability to manipulate graphics objects in near real time. Today there are still only a small number of vendors who can deliver this capability in a workstation environment.

NAS users are for the most part studying problems in CFD. This discipline attempts to solve the differential equations of fluid flow numerically to understand flows over realistic three-dimensional objects. Historically, fluid dynamicists have studied flows with photographs using a wide array of techniques to visually enhance salient features of the flow. The use of computer graphics is a particularly natural step for the CFD scientist. The supercomputer provides him with a "wind tunnel" to perform his numerical simulations; workstations with dynamic graphics capability provide him with numerical flow visualization. Such applications are well described by Buning and Steger, Reference 2.

### 8.3. DESCRIPTION OF PRESENT WORKSTATION HARDWARE

A block diagram of the NAS workstation is shown in Figure. 2. Silicon Graphics Inc. refers to it as the IRIS (Integrated Raster Imaging System) 2500 Turbo workstation. The Motorola 68020 chip and the Weitek chip set give the workstation a floating-point capability greater than that of a VAX 11/780 with floating-point assist.

The unique feature of the IRIS is the Geometry Engine (TM). It consists of 12 VLSI chips that perform the floating-point calculations needed to transform, project, and clip geometrical data for display on a CRT. The IRIS performs transformations at a speed that requires Geometry Engine Performance in excess of 10 MFLOPS. Some rendering of the graphics, including z-buffering (hidden surface) and smooth shading, is also done in hardware. The 32-bit planes of the 1-kbit by 1-kbit resolution display memory can be configured in several ways. For animation two buffers of 12 planes each are used in a "double-buffer" mode. Full RGB color is obtained by using a single 24-bit plane buffer. Hidden surface removal is provided by using 16 planes for a z-buffer and 16 planes for color.

In addition to the usual compliment of RS232 and Ethernet ports, the IRIS also supports a HYPERchannel interface. This is simply a multibus board made by IKON (Seattle) that emulates a DR11W, an interface that permits connection to a Network Systems Corporation A400 HYPERchannel adapter. This interface permits the IRIS to have direct communications to the supercomputer. A four-trunk A400 costs about \$80K and can support four IRISs. The connect cost per IRIS is thus \$22K, including the IKON interface.

### 8.4. THE NETWORK ENVIRONMENT

There are three principal networks associated with the NAS system. The local area network based on Ethernet links all computers with the exception of the supercomputer. Cray will no doubt make Ethernet access available. The lack of an Ethernet interface to the Cray and the requirement for direct access of all computers to the supercomputer requires the use of some facility like the Network System Corporation's HYPERchannel product. Both of these local area networks connect computers located within the NAS facility as well as to computers in other buildings at Ames. These campus-wide connections are made via fiber optics, repeaters, and bridge connections for both Ethernet and HYPERchannel.

The long-haul or wide area network, NASnet, provides Ethernet access for over 70% of NAS users. This network connects the local NAS Ethernet to remote site Ethernets via Vitalink TransLAN communication bridges. The necessary terrestrial communication links are provided by NASA's newly implemented Program Support Communication Network (PSCN). This long-haul network has been in a prototype mode for over two years. In particular, IRIS workstations located at LaRC and LeRC have Ethernet access to the NAS network.

All the computers at NAS run TCP/IP communication protocols. In addition, Berkeley-style networking commands are supported on all systems. This greatly facilitates the addition of new computers to the network. The use of internet protocols provides remote users transparent access to the Cray-2 via Ethernet-HYPERchannel gateways.

## 8.5. THE CRAY-2 AS A GRAPHICS CO-PROCESSOR

In the foregoing sections, the tools for large-scale scientific computing have been described. These are the supercomputer, the graphics workstation, and the communications linking them. An application program known as RIP--remote interactive particle tracer--was written at Ames to use these tools to increase the productivity of CFD users. The UNIX and communications system issues associated with RIP are discussed by Choi and Levit (Reference 3); two specific applications, RIP and PLOT3D, are described by Rogers et al. in Reference 4.

Figure 3 illustrates the RIP application. It involves two processes, one on the IRIS and one on the Cray-2. These processes communicate over Ethernet or HYPERchannel using TCP/IP protocols. The process on the IRIS controls a graphics database for the object under study such as a space shuttle or new fighter design. The IRIS can rotate and zoom through this database in near real time, completely independent of the Cray. The IRIS also provides the principal interface to the user in this application. The process on the Cray-2 controls the solution database for the object under study. This database is the result of solving partial differential equations that describe fluid flow. It is typically 50 Mbyte in size and requires 1 to 20 hours of Cray time to calculate.

The CFD scientist uses RIP by indicating with the mouse a point on or near the object where he would like to release a test particle. The solution database has the information to show him how the test particle flows past the object. When he clicks the mouse, the location of the point is sent to the Cray. It in turn interpolates through the solution database to figure out where the particle will flow. This technique is referred to as particle path tracing. The Cray-2 is very useful for finding particle paths. The search procedure is CPU intensive, and the database is 50 Mbyte or larger. While particle tracing can be carried out on the workstation, the process would hardly be considered interactive. In roughly 1/10 of a second, the Cray returns to the IRIS a series of perhaps 400 short vectors, which geometrically define the particle trajectory.

In a matter of minutes, the CFD user can define and build up a visualization of the flow field over the object of interest. At any point he can rotate the object and the traces to study the flow from different orientations. In this way, he can correlate what happens in one area of flow with another. RIP provides him an interactive tool with which to visualize and explore the results of his flow field calculation. Since RIP sends display list information, not image data, to the IRIS, all viewing and manipulation of geometry and traces are done independently of the supercomputer.

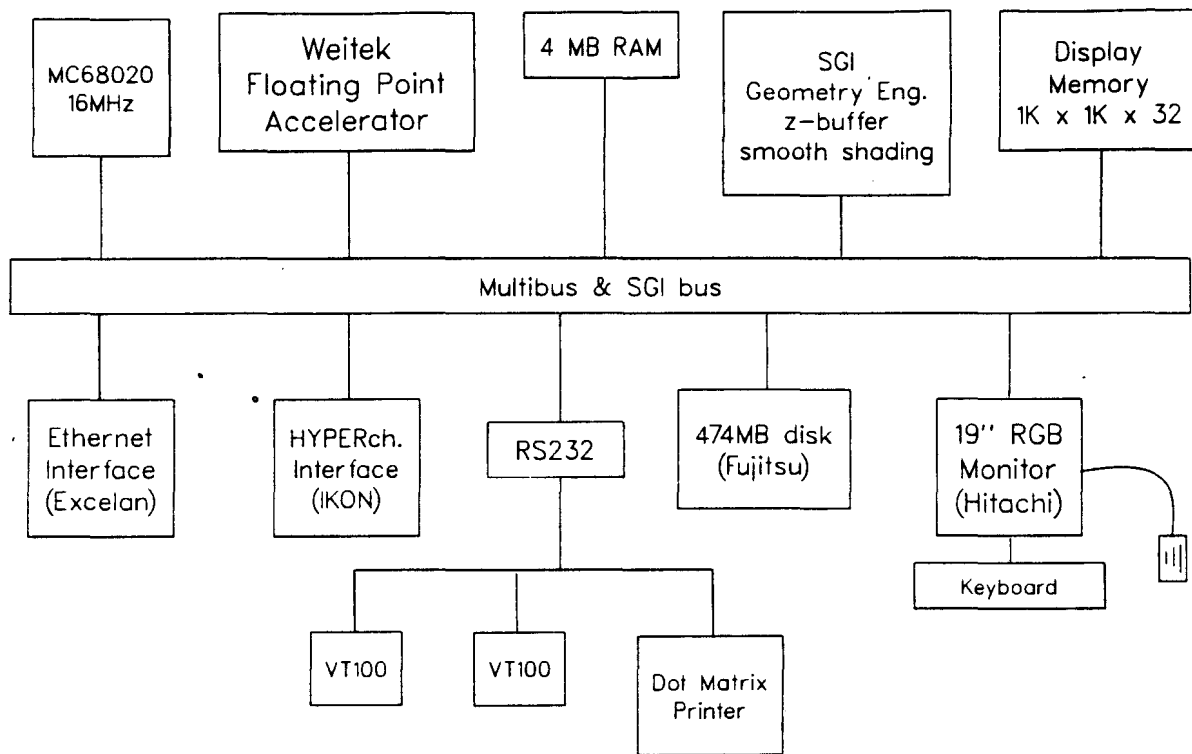


Figure 2. Block diagram of the NAS graphics workstation.

### rip: A Paradigm for Visualization in CFD




Figure 3. An example of distributed scientific graphics.

A video tape illustrating how RIP works is available upon written request. RIP can be effectively used by remote users since it does not send huge amounts of data to the user's workstation. One trace is typically made of about 400 vectors. The tip of each vector is defined by 12 byte (x,y,z). Another 3 to 4 byte are needed for control and color information. At 15 or so bytes per vector, one trace amounts to about 6 kbyte of data to be transmitted. This takes about 1 s at 56 kbyte/s and 1/30 s at 1.5 Mbyte/s. For most users, an interactive response corresponds to 0.25 s or faster.

The problem with particle traces is that they only reveal a fraction of information available in the simulation. In two dimensions, CFD scientists often use contour plots of density or Mach number to gain a more detailed insight into the nature of the flow field. Such plots can be thought of as being made of many traces, typically on the order of 40 to 100 traces. In three dimensions, users now stack two-dimensional contour plots. In time, more sophisticated visualization techniques will be developed. For now one can estimate that the equivalent of about 1000 traces will be involved. Roughly speaking then, CFD users would like to view their data with single graphics frames corresponding to about 1 Mbyte of graphics data. Such views would require 2.5 minutes to transmit at 56 kbyte/s or about 5 s at 1.5 Mbyte/s. These image densities (1 Mbyte) are much more typical of what the CFD scientist would like to study interactively. In this regard, RIP is a compromise and a fortunate situation. The CFD scientist can get some very useful information from the particle traces. This technique may very well be of little or no use to the visualization of other physical problems.

The visualization requirements discussed so far correspond to steady-state flows where the flow field is independent of time. There is no end to the types of unsteady flows that the CFD scientist would like to study. As a dramatic example of such a flow, consider the problem of store separation from a highly maneuverable fighter. In many cases, new store designs or new tactics lead to the store destroying a portion of the aircraft wing. The study of such phenomena is experimentally very expensive. The time-dependent nature of this flow is clear. To visualize it would require hundreds of time steps, with each frame requiring 1 Mbyte or so of data as discussed above. The motion of various components of this flow relative to one another is a critical feature of such a problem. This requires animation on the order of 10 frames per second. This sort of application thus requires data rates on the order of 10 Mbyte/s.

As in other applications, the CFD user also wants the ability to rotate, zoom, and pan through these images either on the fly or in a temporarily paused state. The work of Winkler et al. at Los Alamos (Reference 5) illustrates the power of animating CFD results at rates of 60 Mbit/s.<sup>5</sup> In this environment, graphics images are stored on magnetic disks and transferred as quickly as possible to a frame buffer, providing animation of about 15 frames per second. This approach, however, does not permit the user to interact with the visualization. New external bus technology, such as that of the Ultra Corporation,<sup>6</sup> gives the promise of streaming data from the supercomputer itself to frame buffers at the rate of 100 Mbyte/s. In this case, interactive control can be gained with appropriate software on the supercomputer. This would constitute the world's most expensive graphics workstation!

## 8.6. CONCLUDING REMARKS

An example of a distributed scientific graphics application known as RIP was discussed. This application uses a graphics workstation for the display and control of geometrical data and a supercomputer for CPU and memory intensive rendering operations. Workstation and supercomputer processes are closely coupled with standard networking software based on TCP/IP protocols. The use of a technique referred to as particle tracing keeps data transmission to a minimum. This in turn makes RIP a viable tool for remote graphics workstation users.



As discussed above, it is evident that applications such as RIP can be quite successful in wide area networks running at T1. Although faster rates are always desirable, a truly robust T1 network would not only handle RIP-style applications, it would also permit sizable solution files to be shipped to the remote user's site for more exhaustive study. The word robust is very important. Scientific users want to move large (50-Mbyte) files routinely to and from their home sites. The impression of many users with regard to today's existing 56-kbyte/s networks is that they work fine for mail but are less than adequate for moving files.

While T1 networks may be suitable for the next few years, there are several factors pushing for greater capabilities. As graphics workstations become more commonplace, scientists will develop new techniques to visualize their science.<sup>2</sup> These techniques will require wide area rates in excess of 1 Mbyte/s. There already exist examples of applications where local rates of over 60 Mbit/s have been achieved.<sup>5</sup> Graphics workstations in the next few years will be enhanced significantly in hardware capabilities. Multibus will be replaced by VME. Processing speed will go from the 1-MIPS regime to 10 MIPS or better. Instead of 100K vector transformations per second, there will be systems that transform 300K polygons per second including z-buffering (hidden surface). The advent of 4-Mbyte RAM chips will bring display memories of 256-bit planes. It is clear that future graphics workstations will have the fast buses and large memories needed to handle intensive supercomputer output. The question is whether future wide area networks will have the hardware and software capability to sustain communications between future supercomputers and graphics workstations.

## 8.7. REFERENCES

1. F. R. Bailey, "Status and Projection of the NAS Program," NASA report TM 88339 (July 1986).
2. P. Buning and J. S. Steger, "Graphics and Flow Visualization in Computational Fluid Dynamics," AIAA-85-1507-CP, AIAA 7th Computational Fluid Dynamics Conference, July 15-17, 1985.
3. D. Choi and C. Levit, "An Implementation of a Distributed Interactive Graphics System for a Supercomputer Environment," NASA Ames Research Center, October, 1986; to be published in the *International Journal of Supercomputing Applications*.
4. S. E. Rogers, P. G. Buning, and F. J. Merritt, "Distributed Interactive Graphics Applications," NASA Ames Research Center, 1986; to be published in the *International Journal of Supercomputing Applications*.
5. K-H. A. Winkler, J. W. Chalmers, S. W. Hodson, R. L. Judd, M. McGowen, and D. E. Tolmie, "The Ultra-Speed Graphics Project of the Advanced Computing Initiative at Los Alamos National Laboratory," Los Alamos National Laboratory (1986).
6. J. N. Perdue, Ultra Corporation, 2140 Bering Drive, San Jose, California (1986).



## 9. NETWORKING TRENDS IN DATA COMMUNICATIONS

---

A. Frederick Fath  
Boeing Computer Services

### Abstract

*This paper addresses the changing communications environment brought about by new technology in both computing and communications and its impact on network design. Advances in microtechnology have enabled a large amount of processing and storage capacity to be available to individual users on a local basis. This in turn has changed the way the users access large databases and computational resources. Likewise, the move to digital transmission and switching systems for both public and private networks supporting all types of communications services is providing options that did not exist in the past. Digital transmission rates far exceeding rates that are supportable on analog transmission lines are now commonplace and affordable. The combination of these advances and the impact on network design is the subject of this paper.*

### 9.1. CLASSIFICATION OF DATA TRANSMISSION TYPES

In order to assess the impacts of technology on network design and on computing architectures, it is necessary to make a distinction between different types of data transmission. The distinction is made on the basis of the overall objective of the transmission type. The first type, designated "Response Time Critical," is that type of transmission where the objective is to minimize network delay. The reason delay is important is that for this type of data transmission it is assumed that people or processes become nonproductive while data are in transit. Thus by minimizing the network delay, the end-user productivity is increased. The classic example of this type of requirement is the interactive access of a remote host computer by a user on a terminal or a call setup packet sent between telephone switches to establish a long distance call. In both cases, if the network is slow in delivering data packets, the user must wait unproductively. The general characteristics for this type are short messages with arrival rates that are a function of either timers or human interactions.

The second type of data transmission, called "Throughput Critical," has the objective of high efficiency. For this type of transmission, one desires the maximum utilization of a communication resource in order to minimize costs or to minimize the time needed to complete the transmission. Examples of this type are file transfer applications or possibly interactive sessions with a large amount of data transmission such as would be needed for real-time graphics. The performance measures include the time and cost of completing the required transmission rather than any packet delay measure. The general characteristics for this type are usually large contiguous messages or files sent between computer devices with no human intervention.

## 9.2. PRESENT NETWORKING APPROACH

Most data networks in existence today share an architectural trait that they rely on a fixed network topology. This means that the networks are composed of network processors (that switch and route data packets) and fixed trunks that carry the transmissions between the processors. Local area networks are an exception as they can share the transmission media among many processors. The major reason for this architecture is that transmission costs were the predominate cost of the network, and because of analog technology, the only method to obtain reasonable bandwidth was with dedicated circuits. Once the step was taken to obtain the dedicated bandwidth, then the issue of how to get reasonable utilization had to be addressed if any type of cost control was desired. The packet network technology allowed the traffic from a wide variety of users to be mixed into the network and this in turn led to the utilization needed. Thus the architectures of most data networks today use packet technology on a fixed topology to carry both response time critical data and throughput critical traffic.

This present data network architecture has some characteristics that cause some inefficiency. First, the issue of mixing the two different types of data on the same network causes some problems. For response time critical data, the trunks should never be designed for greater than 50% to 60% utilization. This ensures that the average wait for a packet to have access to a trunk circuit is very low. Further, the packet size for response time critical traffic should be small so that forwarding at network nodes and queuing delays due to other traffic are minimized. Throughput critical traffic on the other hand should maintain as close to 100% trunk utilization as possible. For efficiency, the packet size should be large to minimize the overhead, and the product of window size and packet size should be high enough to allow high efficiency even with acknowledgment delays due to buffering or satellite circuits. There are possible protocols that will lessen but not eliminate the problems of mixing the two types of traffic. Queuing priorities, transmission aborts, dynamic flow control, and other routing mechanisms could be implemented as a compromise to mix the traffic. Few existing network implementations are very successful at efficiently carrying the mix.

The last issue on the present network approach is that of gateways. The fixed topology architecture generally means that users or systems connected to one network will have to use some gateway to other networks if the desired destination of the traffic is not on the same network. Sometimes, it may be necessary to tandem through multiple networks in order to reach the desired destination. For this to happen, it is critical that standard protocols and standard addressing be selected, otherwise the gateway function will tend to be very complex. While the gateway approach to reach destinations not on the originating network has been successful, the network design and optimization task is quite difficult. Traffic engineering for networks with tandem type traffic and alternate routes outside the network usually results in low utilization of transmission resources if performance levels are to be maintained.

Figure 1 is a representation of the architecture of the present network approach. It shows an end-user system connected to a local network with a gateway to a global network. Assuming that the destination system was yet another network away, a gateway between global networks is shown. At each gateway and at each node within the global networks, the packets must be buffered and queued for the next transmission. Further, even with alternate routing, the availability of gateways and global network nodes determines the overall reliability. The more gateways and nodes that a packet needs to traverse, the more delay and chance for failure.

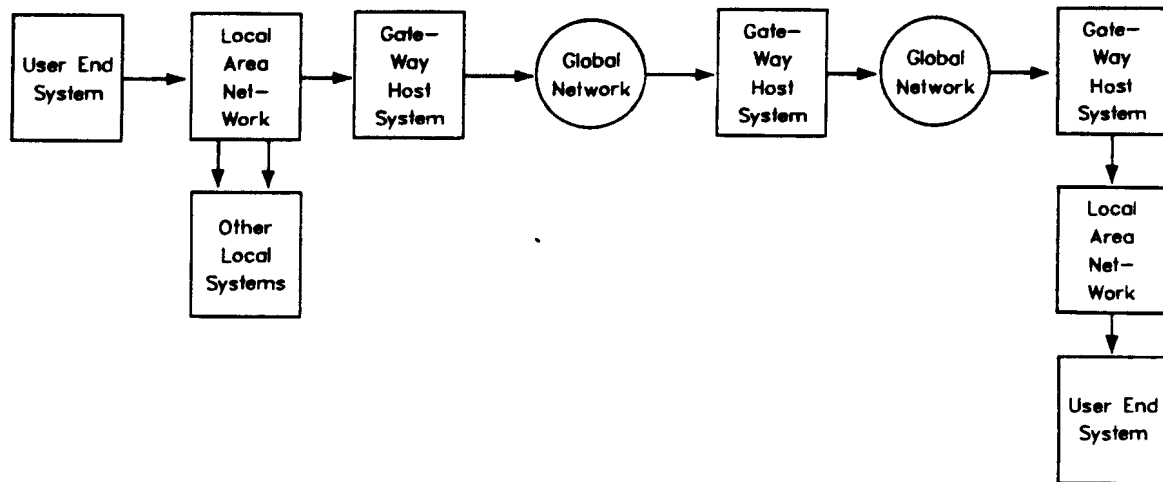


Figure 1. Present network approach.

### 9.3. EMERGING NETWORKING APPROACH

The networking approach that seems to be emerging with the advances in both computing systems and communication technology is one that separates the response time critical traffic and the throughput critical traffic. The characteristics of each type of traffic are taken into account in the selection of the network approach for each. Response time critical traffic with its short statistically distributed messages seems a natural for current packet switched networks. The changes that will be seen in the future center on the ability of the network processors to use dynamic bandwidth for trunking. In periods of low traffic, low bandwidth would be used. When traffic levels increase, more bandwidth could be added. Processors may very well evolve that would be capable of not only adding bandwidth but of creating new direct trunks to nodes as the traffic patterns changed.

The changes in communication technology that would allow the dynamic addition of bandwidth or the creation of new packet network trunks are essentially moves towards total digital switching and transmission. As the dependence on voice grade analog circuits is decreased, the ability to provide higher bandwidth digital paths increases. If network designers have the ability to use circuit switched digital paths in an efficient manner, the opportunities for optimized networks are enhanced. With digital circuit switching, throughput critical traffic can be carried directly from source to destination with circuit utilization at close to 100% with proper protocols. If costs are related to usage and there are multiple destinations, then circuit switching will generally minimize the network costs for throughput critical data. If there is a high concentration of traffic to one location, dedicated facilities at a fixed cost may be required.

Figure 2 illustrates the approach using circuit switching. Once the circuit is established, the impact is the same as the physical interface being extended between the two end-user systems. Depending on the communication media used within the circuit switch network, the delay could be significant. If

satellite technology is used because of its cost advantage, the bit delay from input to output can be around 250 ms. If multiple hops are encountered, as sometimes is needed for international service, the delay will increase. Thus for throughput critical traffic, it is necessary to use protocols that have the product of window size and block size great enough to keep the link busy.




Figure 2. Physical layer switched network approach.

There are advantages of the circuit switched approach for throughput critical data besides cost. Since circuit switching is essentially a physical layer interface and bit transparent, as long as the end systems agree on a protocol suite the connection can be made. There is no requirement for intermediate protocols or gateway functions. While protocol standards are encouraged, the network does not require them. Thus even vendor-specific protocols could be used when appropriate between similar end systems. Further, protocols appropriate for direct connect through such media as satellites could be chosen without impacting other network functions requiring packet level routing. Since circuit switching does no intermediate buffering and the full bandwidth is available, the time to transmit a file once the connection is made is minimized. The network does not require any intermediate host or gateway availability, which makes operations more simple. Finally, the characteristics of throughput critical data follow that of other circuit switched services such as voice, facsimile, and teleconferencing, meaning that integrated services can be combined on network resources to gain the economy desired.

If one follows the separation of traffic types onto different classes of network, it will mean that the end-user systems will have to have multiple network access. Since every user will likely have need for various functions in both the response time critical and throughput critical arenas, it will be necessary for the system to support both types of network connections. Further, there will be local functions supported on various types of servers that will, in all likelihood, need a local area network connection. The result of this is that systems designers will need to address the different types of network connections to obtain desired performance and cost objectives. Figure 3 is a schematic of this structure.




Figure 3. Optimal end user system structure.

#### 9.4. TRENDS IN CARRIER SERVICES

One of the most talked about directions in the public carrier arena is that of ISDN or Integrated Services Digital Networks. This concept of digital circuit switched and packet switched service that allows multiple communication requirements to be satisfied over a standard interface is supported by virtually every carrier both domestic and international. International standards committees have begun setting standards for the user interface at essentially two levels. The basic rate interface would provide two 64-kbit/s circuit switched channels (called B channels) used for voice and data and one 16-kbit/s packet switched channel (called the D channel) for signaling, control, and other packet data services. These logical channels would share a physical channel connecting the interface to the carrier network. A higher speed primary rate interface follows the same structure but with 23 B channels and a 64-kbit/s D channel. Carriers are planning to offer 64-, 386-, and 1472-kbit/s circuit switched service as well as the packet switched service. If one needed higher speed packet service, one could use a circuit switched service to access a higher speed packet switch.

Given that market demand and tariff pricing will make the ISDN an economic approach, then again the systems will need to be designed to take advantage of this network type of network service. This is exactly what was suggested in the previous section.

#### 9.5. NETWORK STANDARDS

Ideally it would be nice if standards existed and were followed for all systems that we use. In reality, it would be a great step forward if even a limited set of standards could be agreed upon for such systems. Most successful standards, however, have evolved from a de facto or dominate supplier system. What results is that standards are chosen for communities of interest as needed and no one standard satisfies all communities. While pursuit of standards is to be encouraged, dependence on standards may be a disaster. It would be nice if everyone in the world spoke English (certainly a recognized standard), but a system depending on it would certainly fail.

The design of networks, especially those carrying integrated services, do depend on standards even though many options exist. Part of the success of the telephone industry has been the acceptance of standards at the low levels and the acceptance of differences at the higher levels. The emergence of

circuit switching for data transmission follows this in that physical layer connectivity will be provided with the functions of the higher layers provided by the end-user systems. This means that community standards for operability can be set without impact on other communities with different requirements and standards. This type of approach extends to other low-level standards such as IEEE 802.3 where multiple high layer protocols can share a transmission resource.

While circuit switching and other low layer protocols allow network sharing, they do little for the higher layers especially when connectivity is needed between different communities of interest. Several approaches exist for solving this problem. First, get everyone to accept a single standard for an entire protocol suite. While this may be possible within a community of interest, it may be quite impractical for larger communities. Second, create a network function that will gateway or convert the protocols to allow access across different standards. Again this may be impractical as the network costs will increase to add this function that many communities may not desire. The third approach is to create, at the end-user system level, an interface with multiple standards. This means a user system would need to emulate or interface with systems of different standards for the set of communities with which connectivity is desired. This means that costs for connectivity are directly related to the number of different community standards that one system must access. If communication is within a single standard, then no additional cost for interface is needed. If a community chooses a standard not supported by many systems, then connectivity may be quite expensive. In this approach, the network would still supply the lower layer communication services but would rely on the end-user systems to correct the higher layer differences.

## 9.6. CONCLUSIONS

The general conclusions that can be drawn as to the trends in data communications are as follows. Packet networks will continue to be the prime method of switching and transmitting response time critical traffic. Circuit switching will emerge as the most cost-effective method for carrying throughput critical traffic. In the near term, common low layer networks will support multiple high-level protocols to similar communities of interest. Since the characteristics of throughput critical traffic are similar to voice and facsimile traffic, circuit switched data will be integrated with voice, facsimile, and even video traffic in an ISDN environment. End-user systems will evolve to support multiple network interfaces for the various types of data transmission. Many end-user systems will emerge that will integrate other user communication services beyond data such as voice, facsimile, and mail. Standards will be chosen within communities and intercommunity connectivity will be accomplished via emulation or protocol matching at the end-user system level.



## 10. INDUSTRY AND TECHNOLOGY TRENDS

---

**Sushil Munshi**  
United Telecommunications Inc.

It is a great privilege to be invited to participate in this Congressional network study concerning requirements and future alternatives for the communications network in support of U.S. research computing. I hope that our contribution will provide useful input with respect to technology and networking trends as envisioned by one of the major communications services providers.

It is clear from the review of the material involved in this study, that traditional views on computer networking requirements are no longer adequate. One of the main reasons for this is the technological advances in computing and networking systems--e.g., microprocessors, memory, local area network workstations, etc. Since these technological advances are expected to continue, a definition of requirements is therefore very critical, particularly of future needs. These requirements can be broken down in various categories: mandatory, optional to become mandatory when technology permits, and desirable (to stimulate technology developments). Range of services, bandwidth control, network management, reliability, security, compatibility with standards, etc., are some of the areas that require careful review.

The obvious description of services generally specifies transport of voice, data, and images. However, it is necessary to define the bit rates (or bandwidth) required and the range covering minimum, say 120 bit/s, to a maximum, for example, of 45 Mbit/s. It is also necessary to define the aggregate demand in order to engineer the transport and processing subsystem adequately. The aggregate bandwidth requirement should further specify if different services are to be provided simultaneously or not. For example, does the end user require video conferencing, voice communication, and data transmission at the same time or on an alternate basis?

Bandwidth control, dynamic or preprogrammed, allows the user to maximize the efficiency of the transport media. The question here is, should it be dynamic (automatic) or preprogrammed with manual override? Since the intelligence is being distributed and is moving away from the central source, an intelligent network front end may well determine the bandwidth to be allocated. As soon as the idea of bandwidth control is entertained, issues of priority and grade of service enter into consideration. Should the bandwidth allocation be based on type of service (priority), tolerance to delay, or other technical constraints, e.g., minimum hops or shortest path?

Network management can be viewed in two different ways.

- One view is the management of assets (transmission) and making effective decisions to meet user needs. Comments on bandwidth control are compatible with this view.
- The second view is the more traditional view of network management encompassing administration, maintenance, update, and usage/performance statistics. The requirement specification should address this second point of view as well.

The security issue can best be addressed by the other distinguished members of the task force who are well versed in the area. I am sure that there are multiple levels of security to limit fraudulent or unauthorized access to the network. However, one should not overlook the privacy issue. Network reliability or availability is of significant importance for the environment proposed. In addition to component or subsystem reliability, the network should accommodate alternate routing or reconfiguration capabilities to maintain high availability. Augmenting the routing schemes can significantly enhance the overall network availability.

Deregulation in the telecommunication industry has created a multiple-vendor environment, particularly in the long-haul communication services. Users, therefore, have a choice of selecting the carrier based on price, quality, or both. It is imperative that interoperability and compatibility among various IECs be defined and specified. Acceptable interface and protocol standards must be mandated by those defining the network alternatives.

Recent technological advances in transmission, switching, and computing systems suggest that the technology is not expected to be a bottleneck for future computing networks. However, the deployment of appropriate technology is essential to support the above statement. Satellites, digital microwave, and optical fiber are available for the transport system. Each one of these media has its place based on economic and technical considerations. Optical systems capable of operating at 1.2 to 1.7 Gbit/s will be available in 1987. Technically, it is feasible to increase the capacity to 4 Gbit/s using the present modulation technique (i.e., no external modulation). Availability of wavelength division multiplexing, coherent detection, etc., is likely to remove any constraints on the bandwidth requirements for the foreseeable future. Japanese firms are introducing fiber based local area networks operating at 45 Mbit/s. Advances in satellite, transoceanic fibers, and digital microwave will definitely satisfy the national, as well as international, networking requirements. Large capacity switching systems (circuit switched) with interface to packet switching and wideband packet switching and, in the future, "burst switching" will satisfy most of the processing requirements. The question here is what technology to deploy: fast circuit, fast packet, wideband packet, or something else? Will the current packet switching systems accommodate the processing requirements of supercomputer networks? These issues can be addressed once the requirements are defined. A number of "boxes" acting as a front-end processor are coming into the market to provide "bandwidth and network control. Digital cross-connect systems will enable network reconfiguration to increase network availability, survivability, and the efficient use of network resources.

It can be stated once more that the needed technology is or will be there to meet the network requirements. However, the most difficult task is expected to be in the transition phase due to the established base of research computing networks. This embedded base requires a strategic approach to build the network of the future and will require economic and/or technical tradeoffs. Having a vision of what the network should look like, one needs to plan how to get there from the current networks in place.

## 11. AT&T PRESENTATION GRAPHICS

---

**Richard Roca**  
AT&T Bell Laboratories

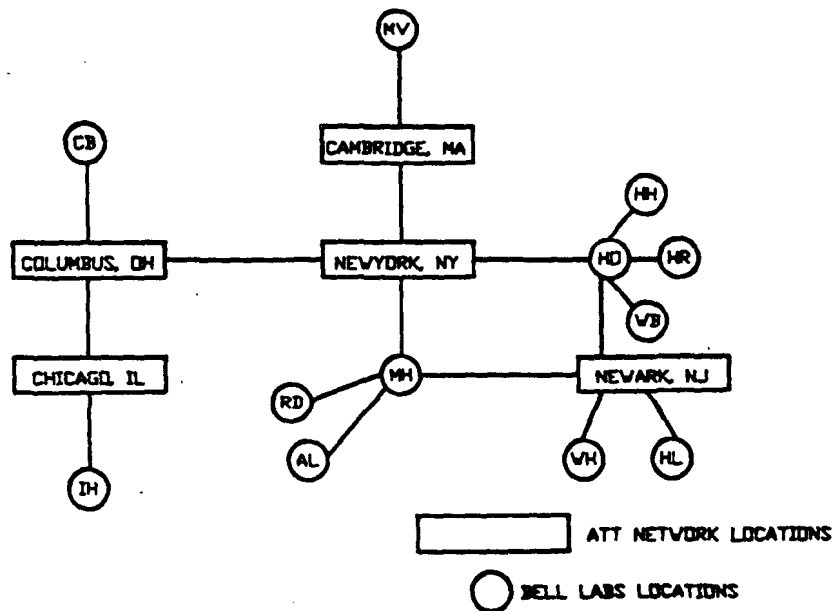
### VIEW OF TRENDS

- DISTRIBUTED INTELLIGENCE DRIVING NETWORK EVOLUTION
  
- KEY INFLUENCING TECHNOLOGIES
  - MICROELECTRONICS
  - PHOTONICS
  - SOFTWARE
  - NETWORKING
  
- INCREASED EMPHASIS ON NETWORKING
  - UNIVERSAL CONNECTIVITY
    - APPLICATIONS
    - TRANSPORT
    - WAN/LAN LINKING
  - STANDARD INTERFACES
  - USER INTERFACES AND DOCUMENTATION
  - NETWORK MANAGEMENT

### PACKET SWITCHED DATA NETWORK

- LOCATIONS THROUGHOUT EASTERN USA
- 50 SWITCHING NODES (DATAKIT VCS TECHNOLOGY)
- 35K END POINTS (60% TERMINAL, 40% HOST)
- UNIX<sup>R</sup> ENVIRONMENT
- TOPOLOGY
  - T1 BACKBONE
  - 56 KBPS INTERNODAL TRUNKS
  - 9600 BPS TERMINAL ACCESS
  - 40-400 KBPS HOST ACCESS
- INTERCONNECTION TO SIMILAR NETWORKS WITHIN AT&T THROUGHOUT USA
- SUPPLEMENTED BY HOST-HOST PRIVATE NETWORK
  - BELL LABS
  - ATT INFORMATION SYSTEMS
  - ATT TECHNOLOGIES
  - ATT/PHILLIPS

## AT&amp;T BELL LABORATORIES T1 DATA NETWORK



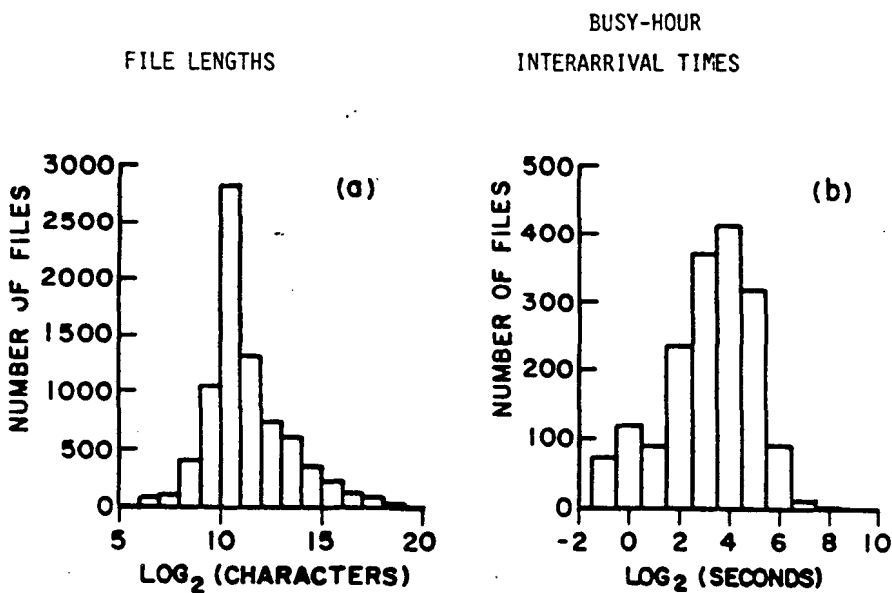
PSDN STATISTICS  
(PER MARSHALL AND MORGAN)

- RESEARCH ENVIRONMENT (1983 AND 1985)
  - COMPUTER AND MATHEMATICAL SCIENTISTS
  - WORD PROCESSORS
  - DUMB TERMINALS AND DISKLESS WORKSTATIONS
- 60% USERS ARE CONNECTED IN BUSY MINUTE
  - 33% CONNECTED TECHNICAL TERMINALS ARE BUSY
  - 66% CONNECTED WORD PROCESSING TERMINALS ARE BUSY
- TRAFFIC
  - 40% TERMINAL - HOST
  - 20% HOST - HOST (TERMINAL LIKE)
  - 40% HOST - HOST (FILE TRANSFER)

BUSY MINUTE TERMINAL TRAFFIC

APPLICATION	TOTAL	TERM TO HOST	HOST TO TERM	MEAN BURST TERM	LENGTH HOST
	(CHARACTERS/SECOND)			(CHARACTERS)	
WORD PROCESSING	10	2.5	7.5	1	40
TECHNICAL (DT)	34	1	33	1	96
TECHNICAL (WS)	72	18	54	8	32


FILE TRANSFER DISTRIBUTIONS  
(PER MARSHALL AND MORGAN)




PSDN - MAJOR ISSUES

- NETWORK MANAGEMENT
  - INVENTORY CHURN
  - TRAFFIC STATISTICS
  - PERFORMANCE MONITORING/ENGINEERING
  - COST ALLOCATION
- INTEROPERABILITY
  - ADDRESSING PLAN
  - DIRECTORY SERVICES
  - GENERAL APPLICATION SERVICES
- FLEXIBILITY
  - EVOLUTION PLANNING FOR NEW TECHNOLOGIES
  - MIGRATION OF PRIVATE LINE NETWORK  
(WAN & LAN SOLUTION)
  - "ODDBALL APPLICATIONS"
- USER SUPPORT
  - DOCUMENTATION
  - INTERFACES


INTEGRATED CIRCUIT TECHNOLOGY



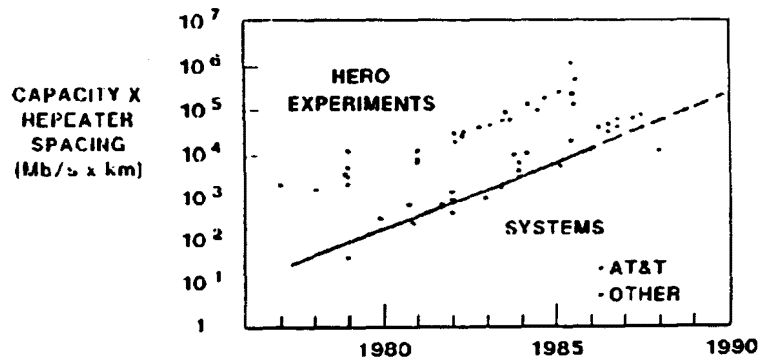
### HARDWARE AND SOFTWARE PRODUCTIVITY




### TREND IN PACKET SWITCHING



### LIGHTWAVE PROGRESS




### ACCUNET T45 DS3 UNIT INVESTMENT

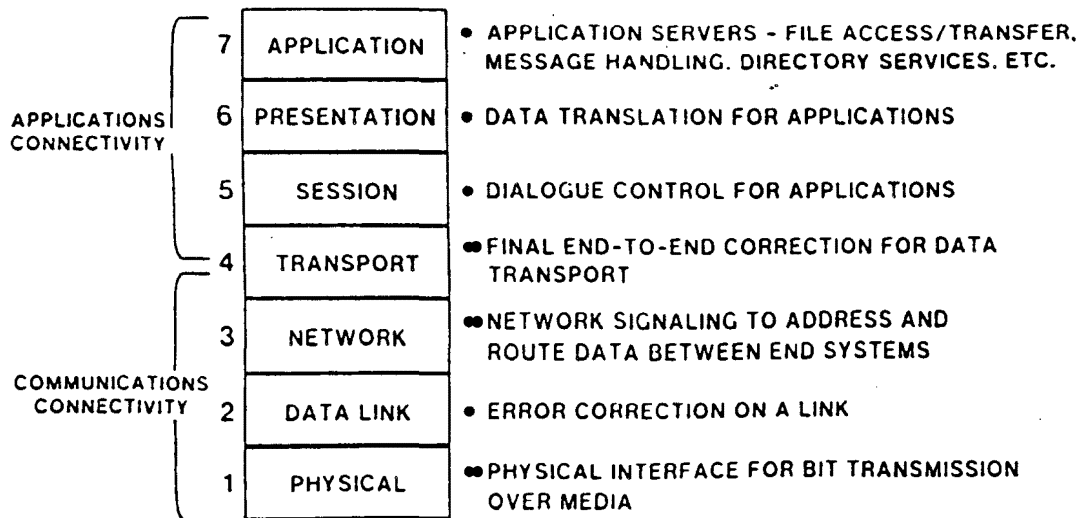




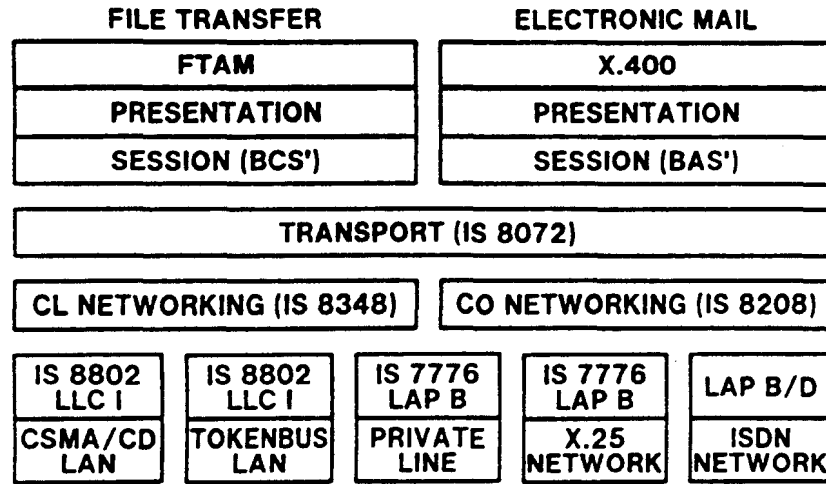
## ACCUNET T1.5 DS1 COMPOSITE UNIT INVESTMENT



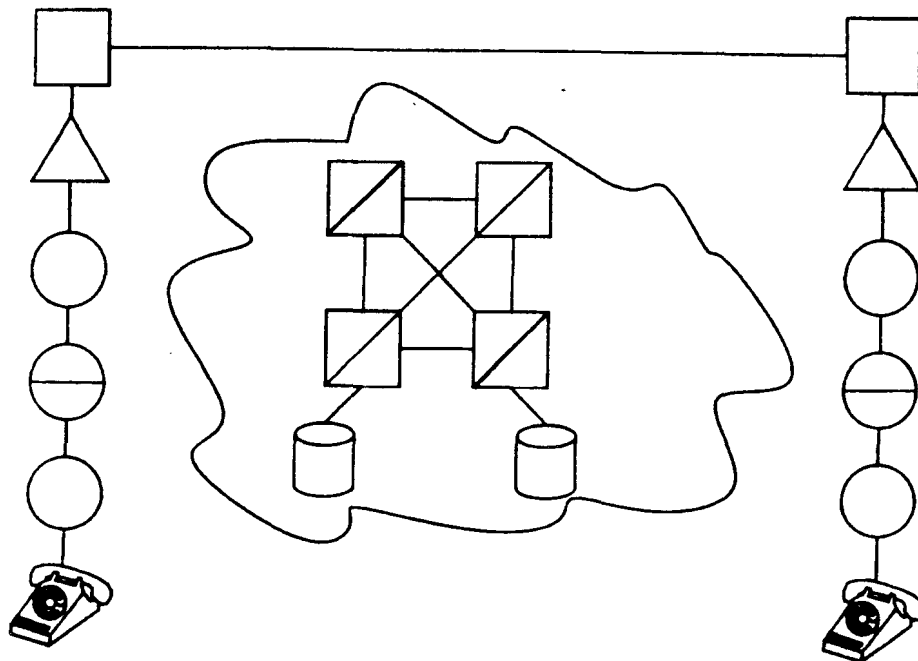
## OPEN SYSTEMS INTERCONNECTION MODEL



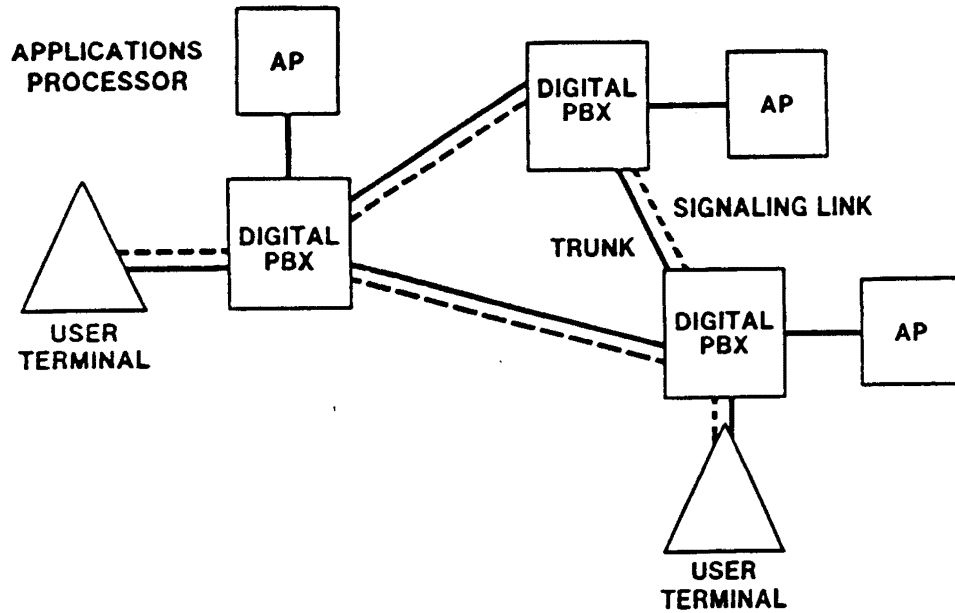
# PROTOCOL SUITES



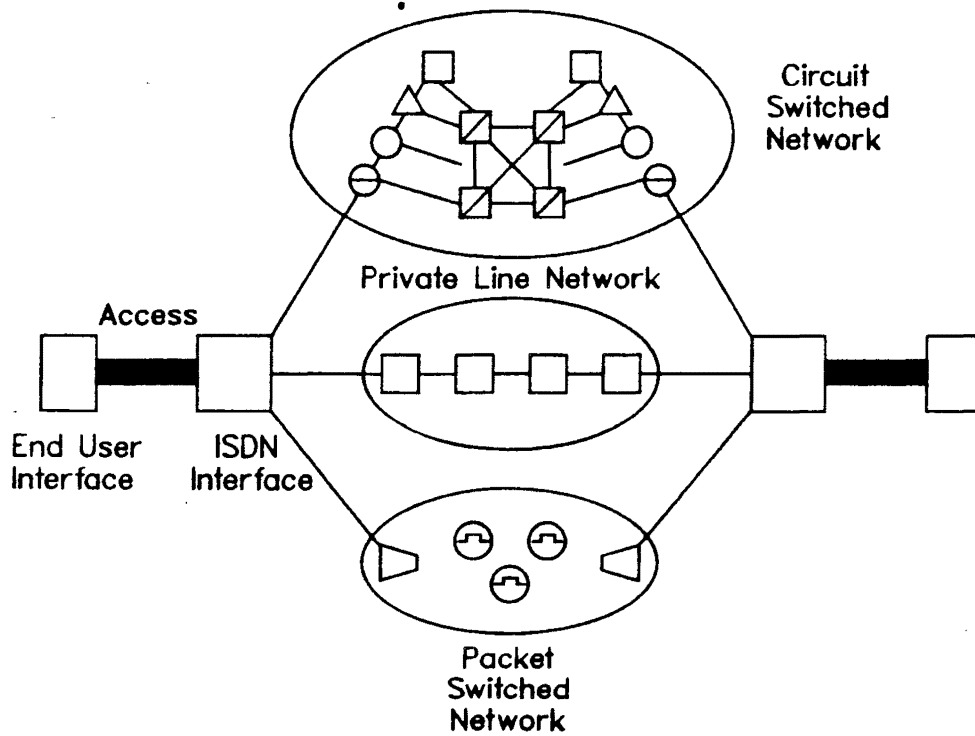
# COMMON CHANNEL SIGNALING




# DISTRIBUTED PBX ARCHITECTURE



# INTEGRATED SERVICES DIGITAL NETWORK




# ISDN



D CHANNEL SIGNALING + PACKET DATA

B CHANNEL CLEAR 64 Kbps DIGITAL ACCESS

# ISDN INTERFACES



### ISDN TRENDS

- BROAD-BAND
- INCREASED FOCUS ON DATA
- GREATER DATA NETWORKING ALTERNATIVES
  - PRIVATE LINE
  - X.25 SWITCHED SERVICES
  - PACKET MODE FRAME RELAYING
- LINK TO NETWORK UTILITIES
  - ROUTING/AUTHORIZATION DATABASES

### ISSUES

- BASIC NATURE OF APPLICATIONS
  - DEGREE OF DISPERSION OF END POINTS
  - DISTRIBUTION OF INTELLIGENCE
  - COMMUNITIES OF INTEREST
  - GEOGRAPHICAL COVERAGE
- DEMOGRAPHICS
- NETWORKING ARCHITECTURE
  - APPLICATIONS SERVICES (MESSAGING, ETC.)
  - TRANSPORT SERVICES - (WAN/LAN INTERWORKING)
  - MANAGEMENT SERVICES
  - EVOLUTION PLANNING
- NEXT STEPS



## 12. CONTEL PRESENTATION GRAPHICS

---

**Patrick V. McGregor**  
Contel Business Networks

Communication Networks for Research Computers  
FCCSET Congressional Networking Conference  
Panel on Networking Requirements and  
Future Alternatives

### **INDUSTRY TRENDS AND COST/CAPACITY FORECAST**

#### **OUTLINE**

##### TRENDS

1. ISDN
2. CCS/SERVICE ROUTING
3. LAN/MAN/WAN/NMS/VATS/SAS
4. TRANSMISSION
5. SWITCHING
6. CONFERENCING

##### FORECASTS

7. CAPACITY/PENETRATION
8. COST/CAPACITY

##### CONCLUSIONS

## TRENDS - ISDN

### - 2B+D

- stronger Centrex
- real potential for residential
- will not displace LAN technology for large intra-site data service

### - 23B+D

- future PBX-LAN/ISDN interface
- will incorporate fast-packet technology
- future computer ISDN interface

## TRENDS - ISDN

- will lead to more efficient and ubiquitous wide area data networking at conventional speeds (<64 kb/s)
- will impact economics of wide area and metro area data networks for interactive services

ISDN WILL BECOME AN EFFECTIVE  
VEHICLE FOR SERVING "1.2-64 Kb/s"  
INTERACTIVE DATA COMM REQUIREMENTS



## TRENDS - CCS/SERVICE ROUTING

- Common Channel Signalling (CCS)  
recognised as big advantage in telephony
- Analogous to separate data/program  
storage in computers (vs von Neumann)
- Trend will be to separate control from  
transport

implies



## TRENDS - CCS/SERVICE ROUTING

1. "Network Control Point" request processing  
centers (vs distributed control)  
note: will reduce vs increase set-up  
delay
2. Service Routing in which different service  
requirements will be served by different  
resource allocations
  - computers and communications
3. Hybrid arrangements ala terrestrial for  
interactive input and satellite for  
high capacity output


## TRENDS - LAN/MAN/ETC

- Local Area Networks (LANs) will continue to expand high capacity intrafacility service
  - 1 mb/s twisted wire pair
  - 10 mb/s shielded pair/coax
  - 100 mb/s fiber optics (today)
  - higher speeds available
- PBXs will be more tightly integrated with LANs to point where integrated systems service full range of requirements
  - integrated with Ethernet today


### LOCAL AREA SYSTEM

- transmission and switching service for a single building or a campus of buildings

**SINGLE BUILDING**



**CAMPUS**



## TRENDS - LAN/MAN/ETC


- Metro Area Networks (MANs) will evolve hybrid microwave/fiber-optic/conventional structures to achieve an expanded single system island.
  - 18 ghz radio "bypass"
  - 23 ghz radio "bypass"
  - fiber optic "dark" channels from local carriers
  - conventional and high speed leased services

LOTS OF HIGHER SPEED CAPABILITY

## METRO AREA SYSTEM

- TRANSMISSION AND SWITCHING NETWORK SERVICE FOR MULTIPLE LOCATIONS IN A METROPOLITAN AREA

- CITY
- COUNTY



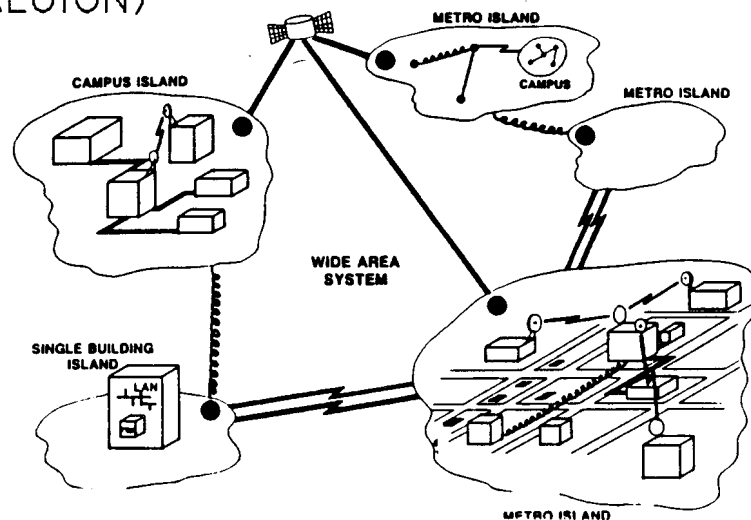
## TRENDS - LAN/MAN/ETC

- Wide Area Networks will suffer pricing disadvantage for higher speed, dedicate channel services (<5 years)
  - lack of fiber optic ubiquity through "last mile"
  - interest in conventional (i.e. 64 kb/s) switching asset use
  - lack of large scale aggregate demand at economical level for ubiquitous network sharing

EXCEPTIONS: METRO ISLANDS, SATELLITE

## WIDE AREA SYSTEM

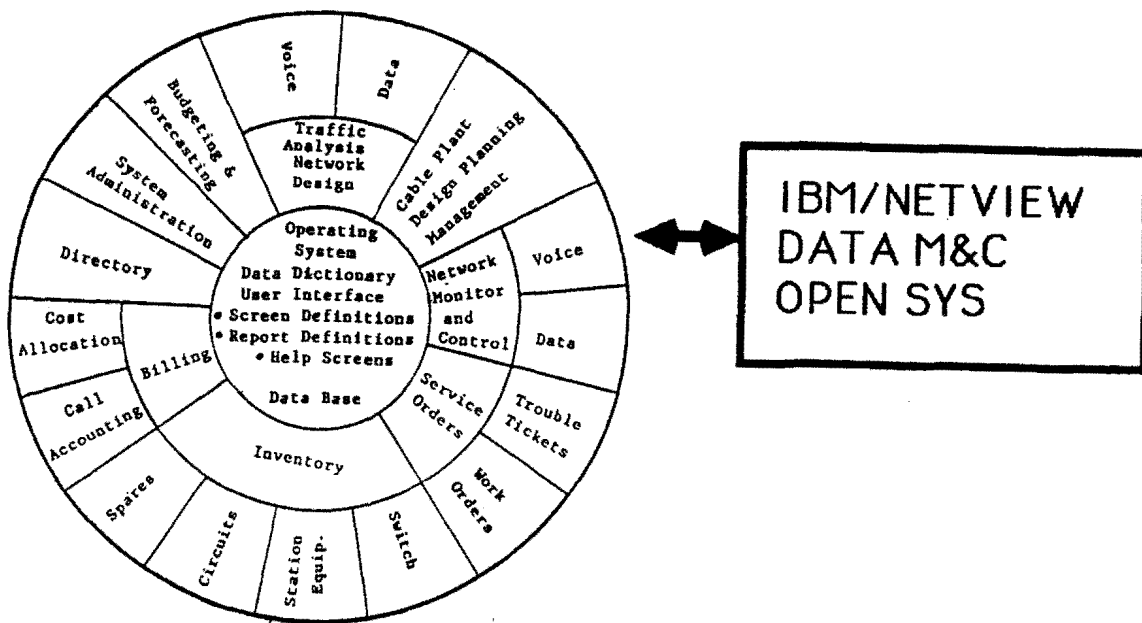
- TRANSMISSION AND SWITCHING NETWORK FOR INTERCONNECTING LOCATIONS DISTRIBUTED AROUND THE COUNTRY (OR REGION)



# TRENDS - LAN/MAN/ETC

- Network Management Systems (NMS) will become increasingly oriented towards full service Network Control Points
  - monitoring and control
  - service routing
  - administration
  - planning and engineering

## network management architecture



## **TRENDS - LAN/MAN/ETC**

- Value Added Telecommunication Services will become increasingly available and integrated
  - Voice Mail
  - Electronic Mail
  - Voice annotated mail
  - Simultaneous and coordinated voice and data
  - Electronic Document Interchange

## **TRENDS - LAN/MAN/ETC**

- Specific Application Services (SAS) will be the primary product thrust of successful market participants
  - value added margins
  - reaching to the user
  - opportunity to eliminate need for wizards

"RICH GET RICHER AND POOR GET POORER"

## **TRENDS - TRANSMISSION**

- MICROWAVE will continue to dominate for most of the intermediate route requirements
- FIBER OPTICS will successfully penetrate heaviest markets and routes, but will not achieve ubiquity (<5 years)
- SATELLITE will offer best avenue for high capacity service to smaller market facilities

## **TRENDS - TRANSMISSION**

- VSATs will advance to provide economical satellite service at 9.6kb/s - 1.5 mb/s from VSAT to VSAT
- Satellite NCPs will permit hybrid structures with terrestrial facilities
- SATELLITE MULTICASTING will become a distinct service advantage

## TRENDS - SWITCHING

- DIGITAL/ISDN SWITCHES will become ubiquitous
- They will serve integrated voice and data
- They will do the data poorly and be behind the power curve from latent services demand
- ASYNCHRONOUS SWITCHING will be required for "ISDN" to fulfill promise - 5 yrs?

## TRENDS - SWITCHING

- "DACS" (T1 switching) will offer some private network relief for high volume users, but will be expensive.
- Satellite switching technologies will make competitive inroads


HIGH SPEED, 'UBIQUITOUS' TERRESTRIAL NETWORKS WILL REQUIRE NEW GENERATION OF SWITCHING TECHNOLOGY (ASYN, FO?)



## TRENDS - CONFERENCING

- teleconferencing will grow
- data conferencing will grow
- vido conferencing will grow the most
  
- effective campus workstation technology available today
- economical at metro area level
- soon at wide area level


BREAKTHROUGH TECHNOLOGIES FOR COLLABORATIVE RESEARCH



# FORECASTS

## VOICE/DATA WORKSTATIONS


(Thousands of units shipped)



# FORECASTS

## VSAT UNITS

(Thousands of units shipped)



# FORECASTS

KB/S	\$/CKT/MONTH			
	100 MILES	500 MILES	1000 MILES	3000 MILES
9.6	476	709	860	1382
56	1988	2995	3861	6156
56VSAT	3100	3100	3100	3100
T1	6246	17282	28083	69132
T1VSAT	15000	15000	15000	15000

# FORECASTS

KB/S	\$/KB-EKT-MILE			
	100 MILES	500 MILES	1000 MILES	3000 MILES
9.6	.49	.15	.090	.048
56	.36	.11	.069	.037
56VSAT	.55	.11	.055	.018
T1	.041	.022	.018	.015
T1VSAT	.097	.019	.0097	.0032

## FORECASTS

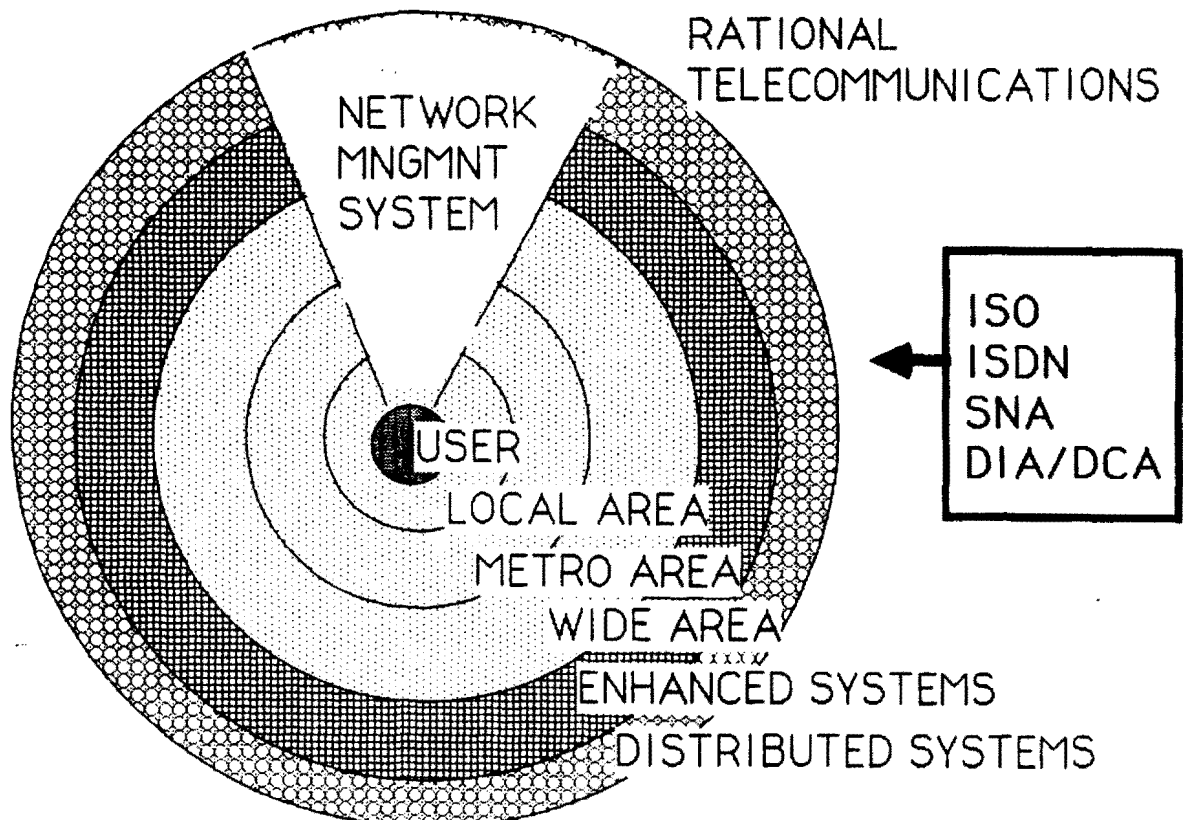
	RELATIVE COST CHANGE					
	86	87	88	89	90	91
9.6	1.0	.95	.91	.88	.86	.85
56	1.0	.92	.85	.79	.75	.73
56VSAT	1.0	.92	.84	.76	.70	.66
T1	1.0	.88	.76	.66	.58	.50
T1VSAT	1.0	.90	.80	.70	.62	.54

## RATIONAL TELECOMMUNICATIONS

COMMUNICATIONS INFRASTRUCTURE THAT

- IS "COST EFFECTIVE"
- PROVIDES PLATFORM FOR APPLICATIONS
- HAS FUNCTIONALITY AND PERFORMANCE  
MATCHED TO USER NEEDS.

**VISIBLE  
INTELLIGIBLE  
MANAGEABLE**



## PRODUCT PLANS

1. 1 mb/s twisted wire pair campus LAN via universal wiring to desk
2. 80 mb/s fiber optic LAN to desk
3. video teleconferencing to desk
4. high capacity, multi-service gateways
5. VSAT family
6. Integrated NMS

## PRODUCT PLANS

7. Value Added Telecommunication Services
8. SPECIFIC APPLICATION SERVICES
9. SKUNK X for breakthrough in keystroke reduction

NETWORK SYSTEM INTEGRATION

## **CONCLUSIONS**

1. To make wide area networks as good as local area networks will require
  - ubiquitous fiber optics
  - new generation switchesneither of which will happen within 5 years.
2. Locations in major markets will be able to form private approximation, but it will be expensive.

## **CONCLUSIONS**

3. Smaller market locations can achieve approximation via satellite.
4. CCS/NCP evolution will make service routing and hybrids more effective and offers significant opportunities for economy and user-friendliness.
5. Satellite multicasting will be significant opportunity for collaborative research.





## NETWORKING REQUIREMENTS



## 13. EFFECT OF DISTRIBUTED COMPUTING ON WIDE AREA NETWORK REQUIREMENTS

---

Dennis Hall, William Johnston, Marge Hutchinson,  
Mendel Rosenblum, and David Robertson  
Lawrence Berkeley Laboratory

### Abstract

*This paper identifies a need to increase wide area network capacity by as much as three orders of magnitude over the next 10 years. These increases are necessary to support new distributed computing products. Such products increase productivity, but are currently available only on local area networks. There is no technical reason for limiting these products to tightly constrained geographical areas, however. They can operate perfectly well over any terrestrial distance provided sufficient bandwidth is available. Such bandwidth is available today with fiber optics. To quantify capacity requirements, network traffic generated by this newer technology is compared with traditional traffic in a local network environment. An extrapolation to wide area networks is made. Speculation about the long-term future of distributed computing technology and its effect on network capacity requirements is offered. It is argued that an increase of network capacity by one order of magnitude is sufficient to accommodate new distributed computing technology on existing wide area networks. Two orders of magnitude are needed to accommodate a fully integrated distributed system such as interactive graphics. Three orders of magnitude are needed to accommodate increases in hardware speed anticipated in the next 5 to 10 years. Availability of a highly integrated, nationwide distributed computing service would significantly increase the competitive edge of the United States in science and computing.*

### 13.1. BACKGROUND

Picture a scientist using a modern, high-performance workstation. One workstation window is opened to a supercomputer located 50 miles away. The supercomputer has been programmed to compute trajectories for a beam of heavy ions from accelerator description parameters stored in a file. The parameter file is displayed on a second workstation window.

Every few seconds, the supercomputer sends a thousand position vectors and a thousand momentum vectors to the workstation. The scientist notices that at each successive step, the momentum envelope is gradually expanding, yet the spatial envelope is holding constant. This indicates heating and means the particles cannot be focused accurately on their target.

The scientist now opens another workstation window. An interactive program is invoked that allows the beam descriptor parameters to be adjusted while satisfying physical constraints expressed as differential equations. A new value for the field gradient in the focusing magnets is established, and the trajectory computation is restarted. The program halts immediately with the message: "Similar values were tried previously: Run 145." The scientist reviews the results of run 145, which have been conveniently displayed by the program, then tries an entirely new value for the magnet setting. The process restarts successfully. This time there is no indication of heating, and the simulated particles proceed to a tightly focused target area.

By using a computer model, an accelerator design flaw has been detected early, and an indication of the cure has been found. The scientist now redirects the graphic display files to a video recorder so that the results can be studied by the full design team at tomorrow's review. Members of the team, some of whom are located at remote sites, are reminded of the upcoming meeting via electronic mail. A few of the images are selected to remind the scientist of the main simulation results. These are printed on a high-quality laser printer located a few steps away, and shared by members of the design team. Additional copies are sent to the remote collaborators.

The preceding scenario is possible using traditional distributed computing technology widely available today. However, these older facilities are more difficult to use and more likely to cause errors than newer products now emerging in the marketplace. These newer products currently operate only in the high-speed, low-latency environments of local area networks. But, their utility is by no means limited to geographically constrained environments. On the contrary, these newer services are useful over any terrestrial distance, if sufficiently high bandwidth is available. In the next section, we review both traditional and more modern distributed computing products.

### 13.2. DISTRIBUTED COMPUTING PRODUCTS--A REVIEW

The following distributed computing facilities have been in common use for at least 15 years, and could all be used in the above scenario. These are the so called *traditional* distributed services.

- *Virtual terminal*: Virtual terminal facilities provide interactive access to remote machines through a computer network.<sup>1</sup> In the example, a virtual terminal facility could be used to access the remote supercomputer.
- *File transfer*: File transfer allows files to be sent from one machine to another.<sup>2</sup> In the example, file transfer could be used to send the particle positions and momentum vectors to the video recorder. It could also be used to send the parameter file from the workstation to the supercomputer and vice versa.
- *Remote job entry*: Remote job entry allows batch jobs to be submitted to remote machines.<sup>3</sup> In the example, remote job entry could be used instead of virtual terminal for controlling the remote supercomputer.
- *Electronic mail*: Electronic mail is the computer analog of ordinary mail.<sup>4</sup> It allows the exchange of electronic "letters." In the example, it could be used to remind the members of the design team of the forthcoming design review.

While the above services are extremely useful, most of them create an unnecessary interface layer between the user and the remote resource. The paradigm of these older services is to provide a remote service through a visible network access mechanism. These were designed in the days of low-performance networks, when network bandwidth was scarce. By making the access mechanism visible, users were made aware of the resource they were consuming. This of course creates extra work, provides opportunities for mistakes, and causes a certain amount of frustration. It has the undesirable effect of reducing productivity.

High-speed local networks have abundant network bandwidth. As a result, more fully integrated distributed services have emerged. In these services, the network is invisible. The following are samples of such services:

- *Distributed printing*: The emergence of high-speed local area networks along with low-cost laser printers has dramatically changed the way printing gets done.<sup>5</sup> It is now economically feasible to allocate high-quality printers to a relatively small group, and to locate these printers in the user's work area. By attaching these to a network, output can be routed to printers thousands of miles

away or a few feet away. Users merely request printing service; the routing is invisible.

- *Network file system:* A network file system makes files available uniformly throughout the network.<sup>6</sup> The machine on which a file resides has no special status. In our example, a network file system could be used as an alternative to shipping the parameter file back and forth between the workstation and the supercomputer. Not only is the time and trouble of shipping files back and forth saved, but more importantly, errors that arise from inadvertently using an outdated copy of the file are eliminated. Users no longer need to be aware of where their files reside. They merely access files in the usual way and location is invisible.
- *Remote procedure call:* A remote procedure call (RPC) is just the ability to call a procedure (or subroutine) on a remote machine.<sup>7</sup> In the above scenario, RPCs might have been used to split the trajectory computation between the workstation and the supercomputer. For example, the program running on the workstation might call several compute intensive subroutines on the supercomputer, which in turn might call graphics subroutines running back on the workstation.
- *Distributed window systems:* Windows provide a point for interaction between user and machine. They increase productivity by allowing users to perform tasks in parallel. A distributed window system permits user level programs to perform complex graphical displays on another machine's window system efficiently.<sup>8</sup> It decouples details such as scaling and positioning from the generating program. Graphical data may easily saturate even the highest capacity networks. Distributed window systems attack this problem by using a high-level graphics description language for communication. This reduces network bandwidth requirements while increasing graphical display functionality. The effect is to make available high-speed, high-quality graphics on machines (supercomputers, for example) whose graphical support system is rudimentary compared to modern workstations. While this technology is still in its infancy, it already promises a major breakthrough in the way scientists interact with supercomputers.

The thesis of this paper is that demand for these newer distributed computing services on wide area networks will increase over the next 5 years. In the next section, we discuss our reasons for this conclusion.

### 13.3. WHY USERS NEED LOCAL AREA NETWORK TECHNOLOGY ON WIDE AREA NETWORKS

Modern, distributed computing products increase productivity. Users accustomed to these products in their local environment will want them in their extended network environment. To illustrate this we compare a traditional implementation of our supercomputer scenario with one using modern distributed computing facilities. We emphasize that both approaches are fully implementable today using off-the-shelf technology. We assume our scientist is using a modern workstation in either case.

In a traditional network environment, our scientist might begin by opening a graphics device emulation window on the workstation. This window would be used just like an ordinary graphics terminal attached directly to the supercomputer. Within the graphics window, a virtual terminal utility would be invoked to access the remote supercomputer. The user would log in to the supercomputer by providing identification (ID) and password information. The remote computation would then be started by providing the names of the program and its data files in a syntax acceptable to the remote supercomputer. The number of steps required to do this would be four: open graphics window, invoke virtual terminal utility, log in to supercomputer, and start program.

In a modern distributed computing facility, our scientist would open an ordinary workstation window, not a special graphics window, and would immediately start the remote computation by providing the names of the program and its data files in a syntax native to the workstation. No special graphics

setup, no login procedure, or other conscious access to the network would be necessary. Moreover, a single syntactic framework, that of the workstation, would be used throughout. The number of steps would be two: open window and start program.

Note that no login step is required in the newer environment. This is normal in a local area network environment where all users work for the same institution and security is not considered a problem. For remote supercomputers, security might well be a problem. To handle this, the supercomputer could request a password whenever the time between accesses exceeds a threshold. The threshold could be chosen to allow users to work unhindered so long as a reasonable degree of interaction is going on. In the above scenario, our scientist would be prompted for a password at most once.

Returning to our comparison, when the scientist notices the beam is heating, the supercomputer computation is stopped in both cases. The next step is to adjust the parameter file. To adjust the parameter file in a traditional environment, our scientist first opens a new workstation window. The parameter file is retrieved from the supercomputer by invoking a file transfer utility. This of course requires providing a login sequence (ID and password) to the supercomputer. The parameter adjustment program is then invoked. Once a new value for the field gradient is established, the scientist returns the file to the supercomputer. The file transfer utility must be reinvoked and the login sequence must be repeated before the file can be returned to the supercomputer. The number of steps would be five: log in, retrieve file, adjust parameters, log in, and return file.

To adjust the parameter file in a modern environment, our scientist simply invokes the parameter adjustment program in the same window used to run the supercomputer computation. Only one step would be necessary: adjust parameters.

If the parameter file is large, and if the network is typical of traditional wide area networks (i.e., 56 kbit/s land lines), transfer time could increase frustration. However, this effect is independent of the utilities used. It is an argument in favor of high-speed wide area networks regardless of the sophistication of facilities.

The next step is to restart the supercomputer computation. Except for syntax, this step is the same in both environments. The program halts immediately with the message: "Similar values were tried previously: Run 145." At this point, the scientist using the traditional environment feels the first real pangs of frustration. Four of the previous five steps must be repeated. (The file needn't be retrieved from the supercomputer initially because it hasn't been changed.) The scientist in the modern environment only repeats one step, the actual parameter adjustment.

This time the program completes successfully. To create the video display in a traditional environment, the scientist must first retrieve the output from the supercomputer and then direct it (as a local file) to the video recorder (two steps). In a modern environment the scientist simply directs the output file to the video recorder (one step). Sending mail to the design team and selecting frames for printing are done the same way in both environments (two steps). However, in a traditional environment, sending copies to remote collaborators would require electronic mail or file transfer rather than a simple print command.

The scientist using the traditional network environment has performed 18 steps while the scientist using the modern environment has performed 7. The effect on productivity is obvious. Run setup time is reduced or eliminated because all resources (remote and local) are accessed uniformly. Errors are less frequent because there are fewer opportunities for their occurrence and because a single command syntax is used. Frustration levels are lower because less time is spent waiting for results,

and because low-level tasks such as shuttling files back and forth between machines have been automated. As scientists become accustomed to these modern facilities in their local environments, demand for similar facilities in wide area networks will increase.

The modern distributed computing environment we have described above is in effect a single, integrated, nationwide "supercomputer." It would be accessed uniformly from anywhere on the network. Its total power would be enormous. Such a facility, available to the national scientific community, would create a technological and scientific environment superior to that of any country in the world. It would help to maintain the nation's competitive advantage in computing for decades to come.

#### **13.4. A COMPARISON OF TRADITIONAL AND MODERN DISTRIBUTED COMPUTING TRAFFIC**

The central theme of this paper is that significant increases in network capacity are needed if local area network technology is to be extended to wide area networks. Our experience in adding such facilities to the Lawrence Berkeley Laboratory local area network is outlined in Section 13.6. We feel this experience provides a forecast of what might occur if modern distributed computing services were extended to wide area networks.

We observe that network file system traffic per host is about an order of magnitude higher than traditional traffic. This is based on observations of diskless workstation traffic. Diskless workstations represent file traffic that can be expected in wide area networks when users must access files from more than one machine (as in the case of our scenario). Therefore, a wide area network that operates comfortably at 56 kbit/s might need a megabit per second (i.e., a T1 channel) to support a network file system or other modern protocols that function at this level of the operating system.

We further observe that more highly integrated services, such as that represented by Sun's memory swapping protocol for diskless workstations (*network disk*), create an order of magnitude higher load than the network file system. Although network disk would not be used on wide area networks because the cost of network bandwidth is much higher than the cost of local disks, it provides a tightly coupled service at a deep level of the operating system. As such it forecasts future distributed computing traffic on local area networks. We conclude that wide area networks operating at T1 speeds might need a 10-Mbit fiber optic link to support traffic from future highly integrated distributed computing services with performance characteristics similar to Sun's network disk protocol. High volume interactive graphics between a supercomputer and a workstation is an example.

So far our analysis has been based on performance of existing distributed computing facilities on existing hardware. The future will certainly bring increases in hardware speeds as well as more highly integrated network software. We think it is reasonable to project a factor of 2 increase every 3 years in available CPU power for the next 10 years. Therefore, the wide area network load can be expected to increase another order of magnitude in 10 years just from faster hardware.

In all, we project an increase of three orders of magnitude in wide area network capacity requirements. In other words, we think that in 10 years scientists could use 100-Mbit links from coast to coast to access a vast array of national scientific computers as a single, integrated "supercomputer." This would significantly increase the competitive edge of the United States in science and computing.

### 13.5. CONCLUSIONS

1. Productivity can be significantly increased by extending modern distributed computing facilities to existing wide area networks.
2. As these facilities become commonplace in local area networks, demand for equivalent services in wide area networks will develop.
3. To accommodate today's network file systems and other highly integrated distributed computing products, a factor of 10 increase in network capacity is needed.
4. To add software products anticipated for 2 to 5 years from now, such as high-volume interactive graphics, an increase of another order of magnitude is projected.
5. To assimilate hardware speeds expected in 5 to 10 years, an increase of yet another order of magnitude is forecast.
6. In all, an increase of three orders of magnitude in wide area network capacity requirements are projected for the next 10 years.
7. Availability of a highly integrated, nationwide distributed computing service would significantly increase the competitive edge of the United States in science and computing.

### 13.6. DISTRIBUTED COMPUTING TRAFFIC AT LAWRENCE BERKELEY LABORATORY

The Lawrence Berkeley Laboratory local area network consists of a single logical Ethernet spanning about half the physical area of the site (several square kilometers). To isolate and minimize traffic, the Ethernet is physically divided into six segments. These are joined by bridging devices with address filtering. This confines network traffic with sources and destinations on the same segment to that segment. A single probe, therefore, can only see network traffic on one such segment. This survey is limited to statistics on two of these segments: CSRLAN, the Computer Science Research segment, and CSLAN, the Computing Services segment. It is further restricted to TCP/IP traffic only. DECnet traffic, the dominant traffic on these networks, is not examined because DECnet currently provides only traditional distributed computing services on Ethernets.

The two parameters used to characterize network load are packet rate and data rate. Tables 1 through 4 summarize the load on CSLAN and CSRLAN in time intervals ranging from 21 to 52 hours. Network traffic is summarized for all the traditional protocols as described in the preceding section. We have included distributed printing in the traditional services because today's wide area networks carry printing traffic, although often it is disguised as file transfer traffic. The left side of the tables characterize traffic as seen by the network. The right side characterizes traffic as seen by an "average" host.

Two lines in the tables are not described in the preceding section. The miscellaneous category covers 18 relatively uninteresting protocols ranging from the internet control message protocol to the time protocol for synchronizing host clocks. The user protocol collects a variety of user developed protocols. Some use Sun's remote procedure call facility. We expect such use to increase as remote procedure calls become easier to use and more widely available.



## Key to Protocol Abbreviations

vt	virtual terminal
ft	file transfer
ml	electronic mail
rje	remote job entry
pr	distributed printing
usr	user defined protocols
misc	miscellaneous
trad	all the above (vt-misc)
nfs	network file system
nd	network disk

Table 1  
Traditional Traffic on CSRLAN--51.99-Hour Sample

service	hosts	kpkts/hr	Mbytes/hr	pkt %	bytes %	pkt/hst/hr	kbytes/hst/hr
vt	124	10.0	0.8	50.2	23.6	161	12
ft	59	0.6	0.2	2.8	7.6	19	8
ml	73	0.5	0.1	2.7	2.4	15	2
rje	18	1.4	1.0	7.1	30.0	157	108
pr	18	0.7	0.3	3.6	8.0	78	28
usr	17	4.0	0.6	20.3	17.5	475	67
misc	106	2.7	0.4	13.3	10.9	50	6
all trad	238	20.0	3.3	100.0	100.0	167	27

Table 2  
Traditional Traffic on CSRLAN--27.52-Hour Sample

service	hosts	kpkts/hr	Mbytes/hr	pkt %	bytes %	pkt/hst/hr	kbytes/hst/hr
vt	93	9.2	0.7	37.0	13.0	197	14
ft	25	0.9	0.4	3.8	7.2	74	30
ml	58	0.7	0.1	2.7	1.6	23	2
rje	14	6.0	3.0	24.1	56.5	854	431
pr	17	1.0	0.3	3.9	5.7	112	35
usr	22	4.2	0.5	17.1	9.9	385	47
misc	70	2.9	0.3	11.5	6.1	81	9
all trad	161	24.8	5.3	100.0	100.0	308	66

service	hosts	kpkts/hr	Mbytes/hr	pkt %	bytes %	pkt/hst/hr	kbytes/hst/hr
vt	79	8.0	0.8	30.0	14.9	203	19
ft	25	0.6	0.2	2.2	3.8	46	15
ml	54	0.4	0.1	1.5	1.0	15	1
rje	22	4.0	1.8	14.9	35.5	364	165
pr	32	9.3	1.7	34.6	34.0	581	108
usr	12	2.7	0.3	9.9	6.3	444	53
misc	66	1.8	0.2	6.8	4.4	55	6
all trad	154	26.9	5.1	100.0	100.0	348	66

service	hosts	kpkts/hr	Mbytes/hr	pkt %	bytes %	pkt/hst/hr	kbytes/hst/hr
vt	38	6.8	0.5	13.9	6.1	360	28
ft	13	0.1	0.0	0.2	0.4	14	4
ml	39	0.4	0.1	0.9	0.6	22	2
rje	14	11.4	5.7	23.2	63.1	1630	807
pr	29	27.5	2.4	55.8	26.2	1895	162
usr	11	1.1	0.1	2.2	1.5	196	24
misc	51	1.9	0.2	3.8	2.2	73	7
all trad	100	49.2	9.0	100.0	100.0	984	179

Table 4 shows a significantly higher value of remote job entry and distributed printing load. These services are the two most heavily used in the traditional set, and use is increasing. Remote job entry traffic derives from the familiar remote shell command in the UNIX environment.<sup>9,10</sup> It allows commands to be executed on remote machines and files to be copied across the network. This latter use is functionally equivalent to file transfer except that no login is required. In the Lawrence Berkeley environment, most remote job entry traffic is generated by disk back-up demons that wake up in the middle of the night to copy UNIX disks onto the central VMS file cluster. The distributed printing service is also popular. It supports the Laboratory's distributed computing service that currently produces about a hundred thousand pages per month of printed output. The percentage columns show relative network traffic on the segment.

The packet and data rates for "average" hosts may be used to estimate the increase in network traffic that would be brought about by adding a protocol to a machine. Note that since all traffic is between two or more hosts, we double the network values to compute hourly rates per host. Clearly, remote job entry and distributed printing have the most significant effect in the traditional set.

The last line summarizes traditional network traffic. The network sees 20-50 thousand packets per hour and 3-9 Mbyte of data per hour around the clock. An average host sees up to a thousand packets per hour and 30-180 kbit of data per hour around the clock. Note that the average traffic rates are

much smaller than the rates for some individual protocols. This is because the denominator in the equation, the number of hosts, includes all machines in the sample. Thus, average traffic is highly biased toward protocols that run on the greatest number of systems. We next show the effect of adding modern distributed services to this environment.

Tables 5 through 8 show traditional network traffic together with two modern distributed services: network file system and network disk. The network file system is described above, and network disk is Sun's network disk protocol (proprietary). Network disk provides a memory swapping service for diskless workstations. It would not be used on wide area networks because the cost of network bandwidth is much higher than the cost of local disks. However, it provides a tightly coupled service at a deep level of the operating system. We expect to see more such highly integrated services emerging in local area networks over the next two to five years. Systems such as Andrew<sup>11</sup> and Mach<sup>12</sup> provide a preview of things to come. Therefore, we feel network disk provides a good predictor for future services that will evolve first in local area networks and then be desired in wide area networks.

service	hosts	kpkts/hr	Mbytes/hr	pkt %	bytes %	pkt/hst/hr	kbytes/hst/hr
trad	238	20.0	3.3	30.9	8.0	167	27
nfs	19	5.0	2.7	7.8	6.6	529	285
nd	16	39.6	34.9	61.3	85.4	4952	4360
total	238	64.6	40.9	100.0	100.0	542	343

service	hosts	kpkts/hr	Mbytes/hr	pkt %	bytes %	pkt/hst/hr	kbytes/hst/hr
trad	161	24.8	5.3	43.9	16.5	308	66
nfs	18	4.7	2.8	8.3	8.5	521	307
nd	15	27.0	24.3	47.8	75.0	3598	3244
total	161	56.5	32.4	100.0	100.0	702	403

service	hosts	kpkts/hr	Mbytes/hr	pkt %	bytes %	pkt/hst/hr	kbytes/hst/hr
trad	154	26.9	5.1	68.1	32.5	348	66
nfs	8	1.0	0.4	2.5	2.5	245	99
nd	9	11.6	10.2	29.4	64.9	2575	2271
total	154	39.4	15.7	100.0	100.0	512	204

service	hosts	kpkts/hr	Mbytes/hr	pkt %	bytes %	pkt/hst/hr	kbytes/hst/hr
trad	100	49.2	9.0	79.9	43.9	984	179
nfs	10	1.6	1.0	2.6	5.0	324	202
nd	8	10.8	10.4	17.5	51.2	2691	2612
total	100	61.6	20.4	100.0	100.0	1232	408

On the CSRLAN network file system, traffic per host is clearly about an order of magnitude higher than traditional traffic, and network disk traffic is an order of magnitude higher than on the network file system. On CSLAN, this doesn't show as clearly because the workstations on CSLAN are primarily used as front ends to VMS machines, and because of the previously mentioned heavy use of remote job entry and distributed printing. Nevertheless, the network disk traffic is a factor of 20 higher than all traditional traffic. We conclude that a wide area network that operates comfortably at 56 kbit/s might need a Mbit per second (i.e., a T1 channel) to support a network file system. Similarly, wide area networks operating at T1 speeds might need a 10-Mbit fiber optic link to support traffic from future highly integrated distributed computing services with performance characteristics similar to Sun's network disk protocol.

### 13.7. REFERENCES

1. "Telnet Protocols," *The ARPANET Protocol Handbook*, Defense Communications Agency, pp. 89-264 (January 1978).
2. "File Transfer Protocols and Standards," *ibid*, pp. 265-332.
3. "Remote Job Entry Protocols," *ibid*, pp. 379-420.
4. "Mail Protocol and Standards," *ibid*, pp. 333-379.
5. W. E. Johnston and D. E. Hall, "Unix Based Distributed Printing in a Diverse Environment," *Summer 1986 USENIX Conference*, Atlanta, Georgia, June 9-13, 1986.
6. "NFS Protocol Specification", *Networking on the Sun*, Part No: 800-1324-03, Revision B of 17, Sun Microsystems, Inc. (February 1986).
7. "RPC Protocol Specification", *ibid*.
8. *NeWS Preliminary Technical Overview*, Sun Microsystems, Inc. (October 1986).
9. "rsh-remote shell," *Unix Programmer's Manual*, 4.2 Berkeley Software Distribution, Virtual VAX-11 Version (March 1984).
10. D. M. Ritchie and K. Thompson, "The UNIX Timesharing System," *CACM*, 17, 7 (July 1974).
11. J. H. Morris, et al, "Andrew: A Distributed Personal Computing Environment." *CACM*, 29, 3 (March 1986).
12. R. Rashid, "Threads of a New System," *Unix Review* (August 1986).

## 14. IMPACT OF DISTRIBUTED FUNCTIONS ON NETWORK REQUIREMENTS

---

Richard G. Zwakenberg and Kirby Fong  
National Magnetic Fusion Energy Center

### Abstract

*This paper considers the current and potential distributed applications that may be used on the network of the National Magnetic Fusion Energy Computer Center. The term "network" is used in its broadest sense to cover not only the reliable transmission of data from one location to another, but also the higher level communications conventions needed to support orderly and efficient development of such applications.*

### 14.1. INTRODUCTION

The National Magnetic Fusion Energy Computer Center (NMFECC) is located at the Lawrence Livermore National Laboratory near Livermore, California. The purpose of the NMFECC is to provide the resources needed to support the large-scale computing needs of the Office of Energy Research in the Department of Energy. The NMFECC maintains and operates a system of supercomputers (currently a CRAY-1A, a CRAY-1S, a CRAY X-MP/22, and a CRAY-2) to supply the major resources. These computing resources are made available to researchers throughout the United States and at selected locations abroad through a data communications network called the MFEnet.

The NMFECC has been exploiting one area of distributed processing for several years: namely, the coupling of the responsiveness of a personal computer with the processing power of a supercomputer to give the users of the NMFECC the benefits of both.

This paper will describe some of the work that has been done in the particular area of distributed processing as described above, some of the plans that are being considered to broaden our capabilities, and finally discuss what the impact and requirements of our plans would be on the national network.

### 14.2. WHY DISTRIBUTED PROCESSING

There are many reasons why a distributed approach to a particular data processing application may be attractive. We have been exploring this area for some time with the primary purpose of providing a more responsive interaction for remote users while maintaining access to the power of several Cray supercomputers.

One does not have to be very perceptive to note the trend toward ever-increasing computer processing capability that an individual may now have available in his office as a personal computer (PC) or workstation. Three trends that can be expected to continue for some time are the following.

- PC resources will continue to become more cost effective; i.e., they will have a growing capability for a fixed cost.
- The cost per unit of data communication will decrease for satellite service and fiber optic transmission. Closely related to this declining cost, the cost per unit of bandwidth will also decrease.

- Supercomputers, by definition, are the fastest commercially available scientific computers, regardless of cost. (However, the realities of the marketplace will probably continue to keep these machines in the \$15M to \$20M range in today's dollars.) The net processing power of such machines will continue to grow, although it seems clear that significant advances will be achieved only with multiprocessor systems with increasing numbers of processors per "mainframe."

Given the above trends, the result is fairly easy to predict. Computer users will continue to enjoy increasing power in their personal computers and increasing bandwidth and lower cost for data communications when connecting to other hosts. However, in spite of this, even when the much prophesied "Cray on your desk" finally arrives, supercomputers will be available with orders of magnitude more capability than the desktop Cray. For many classes of scientific research, access to such supercomputer resources will remain essential.

Although all three elements of this environment are steadily improving, perhaps nearly constantly relative to each other, the absolute power and data communications bandwidth available at the user's fingertips will accentuate the need for effective distributed processing, and may continually change the "right" place to split a distributed application.

To see distributed processing in the NMFEC's context, one should realize that the overall goal is to provide convenient, interactive, supercomputer services. This goal is partially fulfilled by using a timesharing operating system on the supercomputers; however, complete success is thwarted by three problems. One is that, under heavy load, no system can provide immediate response to every user. A second is that supercomputers typically are not designed to respond to every keystroke but are line oriented so that a desirable service such as screen editing is difficult to implement. A third is that the data communications network between the supercomputer and the user may include delays (e.g., due to satellite transmission) that make immediate response to the user impossible. We therefore wish to use the power of a personal workstation (hence distributed processing) to give us a more responsive, interactive interface to supercomputers. The role of distributed computing, or distributing a function, is to take a function and partition it, splitting the interactivity from the calculation and moving the highly interactive portion to a more interactive computer, namely the user's personal workstation. Making workstations more useful is not the objective; it is a byproduct of our overall goal of making supercomputing easier, more interactive, more responsive, and more convenient.

At the NMFEC, we plan to support, in the short range, specific distributed computing applications and, in the long range, user- or vendor-distributed computing applications based on industry standards. Planning for a distributed computing environment requires the development of basic software building blocks and research into some important questions.

### 14.3. SPECIFIC DISTRIBUTED APPLICATIONS

In the short range, we are developing specific distributed applications to achieve our goal of a more responsive supercomputing service. These applications center around remote terminal access, handling of keystrokes, cursors, etc., and the transferring of files and blocks of information between the supercomputer and the workstation. The first of these distributed applications we have developed is a distributed screen editor.

This distributed editor has been partitioned into two pieces. The file management and calculation intensive portion runs on the supercomputer, and the interactive portion (keystroke or screen handling) runs on the workstation, an IBM PC. The supercomputer does such things as searching for patterns, replacing patterns, and keeping the entire file up to date, while the PC handles the character changing.

deletion, and insertion. This splitting of the work between supercomputer and PC makes for a highly interactive text editor, a text editor that can work on the large files that are often created by users of supercomputers.

To make more effective use of the supercomputer and the PC, we have developed a windowing capability on the PC that allows a user to interact with many programs at once. The user can view the output of a program in one window, be editing a file in another window, and be watching the results of, for example, a compilation in another window. This output window capability is not really a distributed application but is an effective way of getting more information into the hand of the user; it has proved to be an excellent productivity enhancement. True distributed computing will come later when we implement multiple input through windows so that the windowing software becomes a program on the workstation that is able to communicate with multiple remote hosts simultaneously.

Another distributed application that we are currently working on is a distributed graphics editor. The graphics file has presumably been produced by an applications program on the supercomputer and then is to be viewed or modified. The graphics editor, like the textual screen editor, is split into two portions, a program running on the supercomputer managing the graphics files and doing the calculation intensive updating, and a PC program that displays the graphics file, allows cursor manipulation, does zooming, panning, etc.

#### 14.4. LONG-RANGE DISTRIBUTED APPLICATIONS

Other distributed applications we plan to develop are described below.

- Electronic blackboard. With this device, two or more users can view the same information in separate (perhaps remote) offices and use their workstations as display devices and graphics input devices. This is to be done in real time as an aid for telephone conversations.
- Remote videocassette recording. In this application, a workstation accepts information via the network for recording on a VCR. There are already VCR controller cards that one can purchase for an IBM PC.
- Menu interface to the CTSS utilities. This development involves either icons or pop-up menus. This distributes the user interface (help or how-to-use package, determining what action the user has selected) to the workstation and leaves only the core of the utility routine on the supercomputer.
- Full input/output windows. These windows are to be capable of addressing multiple remote hosts simultaneously.

We also expect vendors and users to develop distributed applications. Users will most likely start by treating their workstations as remote graphics stations, which are able to accept, process, and display scientific data. That is, the workstation is able to accept data in a compact form from which to generate a graphical display. Vendors of workstations may look at supercomputers as back-end machines and offer software to integrate the supercomputer functions into the workstation's operating environment. That is, a request is executed on the workstation, if possible; otherwise it is sent to the supercomputer. There is already considerable work on distributed databases involving multiple hosts and multiple simultaneous users (not just one user and two hosts). Scientific databases (say of bibliographic abstracts or experimental data) can be maintained and accessed with distributed computing--the database could reside and be maintained on the supercomputer and querying programs could be run on workstations. All of these applications presume there will be some way for two or more machines to talk to each other and transfer information to each other. These very basic capabilities to support distributed applications ought to be available in a generally useful form as building blocks.

A specific user application will be run-time monitoring of, for example, physics or hydrodynamic codes. The user will want to write his or her own program that permits interactive observation of the code's progress and run-time intervention if necessary. We anticipate an interactive program running on the supercomputer that sends output to the workstation for the latter to plot while the supercomputer program proceeds. The user will be able to generate various displays on the workstation that then suggest how to modify the course of computation on the supercomputer. Various pieces of software (probably in the form of user callable library subroutines) must be available to the user for constructing such a distributed application.

#### **14.5. BUILDING BLOCKS AND TOOLS FOR DISTRIBUTED APPLICATIONS**

The development of specific, distributed utility routines, while very valuable to the user community in general, is still not sufficient. It is a necessary first step in understanding how parts of a distributed program must communicate with each other and in acquiring a perspective with which to evaluate proposed distributed application protocols. Protocols are simply conventions to which many programs agree to adhere.

Protocols are often made available to programs in the form of libraries of procedures that can be called. In this way, strict adherence to a convention is embedded in a library, and the library in turn presents a simple interface to the program. There are currently a number of proposed protocols for such diverse activities as

- one host computer invoking a program on another host,
- one host accessing a remote data file,
- exchanging numeric or other forms of data, and
- windowing.

None of these is close to becoming a standard.

A standard is a convention that has been officially adopted by the American National Standards Institute. One current standard specifies the Graphical Kernel System (GKS), a device-independent applications interface to graphics. The NMFEC will implement a GKS that distributes some of the graphics workstation functions to the workstation itself rather than perform them all on the supercomputer. Another very new standard is for a device-independent graphics file, the Computer Graphics Metafile (CGM). This file format and its extension, which is yet to be standardized, should become the medium of exchange for graphics information throughout the network. The routines for reading, which are yet to be standardized, should become the medium of exchange for graphics information throughout the network. The routines for reading and writing this file should be added to GKS to insulate the application programs from the details.

A network, strictly speaking, is made up of the lowest three layers (physical, link, and network) of the International Standards Organization's seven-layer Open Systems Interconnect model. Beyond this are the transport, session, presentation, and application layers with which we must become concerned to achieve the services we are striving to provide. As an example of how these higher layers come into play, we have heretofore permitted the supercomputer to view the personal workstation as a terminal. In the long run, the network must be able to handle the personal workstation as though it were a host computer so that there will be a more symmetric relationship between the supercomputer and the personal workstation. This particular feature will permit the workstation to connect to multiple remote hosts simultaneously or to a centralized data base manager to talk to multiple users simultaneously. Furthermore, we must support the notion of a dialup, temporary, remote host since we cannot expect



all of our users to be permanently connected to MFEnet II. Once the user's workstation is connected to the network, further enhancements may be needed to the protocol to use alternate network paths effectively to select the paths with the most appropriate transmission characteristics for a particular application. Again, whatever protocol is eventually used should be available to applications programs in the form of libraries of callable procedures.

In addition to building blocks for distributed applications, there must be tools to partition existing applications, tools to analyze the performance of a distributed application, and tools to debug distributed applications. Very little of this type of software now exists. A debugger for distributed applications would itself probably be a special, distributed utility program, having parts that run simultaneously on the same hosts as the distributed application it is trying to debug. Partitioning tools must take into account the power of each host and the transmission characteristics between them to analyze an application program for effective partitioning. Such tools will likely assume there is no other competition for network or computer resources when performing a partition; however, in reality, other users and their distributed programs will keep computers and the network busy so that there will be delays or bottlenecks for which such tools cannot account. Therefore, performance measurement tools are needed to determine whether a distributed application is working satisfactorily in a busy environment.

## 14.6. NETWORKING REQUIREMENTS

Current and future distributed applications will make demands on all levels of the network. Item 1 below concerns the network proper; the other items deal with the transport layer and higher.

### 1. Bandwidth, Delays, and Alternate Routing

The bandwidth impact of distributed applications seems to be somewhat of a "chicken and egg" problem. Development of applications that require higher bandwidth than available will not generally be undertaken until the bandwidth is available. On the other hand, managers of networks are reluctant to upgrade bandwidth until the need is demonstrated. Distributed applications will generally be designed around available bandwidth, or as in our early efforts, will be specifically designed to minimize their bandwidth requirements. This dilemma is somewhat inherent, since one of the major parameters that must be considered when distributing an application is how much time will be required to exchange data between the computers involved, i.e., how much data must be exchanged and what is the effective bandwidth and delay. We hope this problem can be addressed partially with protocols that are adaptive; that is, they can adapt dynamically to the particular transmission characteristics involved.

However, given that both the supercomputer and the workstation environment will continue to grow rapidly more powerful, it is fairly safe to expect that the total data traffic caused by distributed applications will also grow. As the local environment becomes more powerful, the user will increasingly resort to new, and often data intensive, local programs for the postprocessing of supercomputer results. Much higher resolution color graphics, with "instant" response and display of local, pseudo-movies, are just a couple of the possibilities being mentioned. The issue of whether high bandwidth, low delay data communications is required is really not in question; the question is primarily an economic one of how much is enough.

One means for the network to respond to the uncertainty of bandwidth requirements is simply to install a great deal of very high bandwidth, low delay capability and hope that the applications will be developed to use it. Unfortunately, this is generally not cost effective since the network must pay for the extra unused capacity until (or if) it ever becomes used.

Additional dilemmas currently exist in the selection of economical communications facilities:

- Terrestrial circuits are attractive due to the fact that they do not suffer from the 0.25-s delay inherent in satellite circuits. However, above 56 kbyte/s, terrestrial circuits become increasingly less cost competitive with satellite circuits.
- Fiber optic links offer the promise of comparable cost, high-bandwidth, low-delay communications. However, currently fiber optic links suffer from limited availability and restricted geographic distribution.
- Satellite circuits offer cost-effective, high-bandwidth data communications capability. Also, with the decreasing cost of earth stations, particularly with the growing availability of Ku band service, a private station can be established at nearly any domestic geographic site desired. Another advantage of satellite communications is that it is inherently "broadcast" in nature; i.e., everyone receives all the transmitted data. This feature can be used to do dynamic assignment of bandwidth. If a receiving site momentarily requires high bandwidth, the required data communications bandwidth can be assigned to it. Since the bandwidth is shared and dynamically assigned, a communications scheme taking advantage of the broadcast capability could prove to be the most cost-effective means of providing very high bandwidth. The major disadvantage of satellite communications is, of course, its 0.25-s delay.

## 2. Access Control

Dialup hosts and hosts in other networks that are interconnected present questions about validation. Who is the user? Is he/she entitled to connect to the network? What are the characteristics of the user's host/workstation/terminal that the network must know? How do we assign a temporary host number or identifier? How can other hosts verify dynamically the validity of a temporary host? What should we do with data being sent to a host that has disconnected? How do we accomplish a reconnection?

## 3. Accounting

The network must account for all traffic in order to charge users and in order to analyze traffic patterns for excess or inadequate capacity on various nodes and links. Due to increasing interconnections with other networks, we can expect traffic to come in through one gateway and out another gateway which is caused by someone who is not a user of the NMFEC. Perhaps some agreement is needed that would be similar to an international postal union, where two interconnected networks agree not to charge each other as long as the traffic is balanced in both directions, and the network of origin charges the sender more for an out-of-network transmission. Should there be such a thing as COD or collect calls for network traffic?

## 4. Reliability

Some distributed applications like videocassette recording require many hours of data transmission. Interruptions may be due to failures in the remote hosts themselves or common carriers (e.g., telephone lines or commercial packet switched networks) as well as the network. This means we need protocols that allow hosts to reestablish connections after some disruption, and this re-establishment must be automatic and not require user intervention.

#### 5. Security, Privacy, and Encryption

If 100% validity checking is not possible (or affordable) or the nature of a data link (e.g., satellite broadcast) permits interception of data, it may be necessary to encrypt data transmissions so that users cannot accidentally or deliberately intercept each others' data.

#### 6. Internetworking

National networks are already interconnected, and this trend will continue. Also, individual sites will connect their local area networks to a national network. Can distributed applications still work when they span two or more networks? That is, will distributed applications require a protocol that may not be able to work through some gateways?

#### 7. Data Compression

Is there an affordable way to compress data in order to reduce the required bandwidth for distributed applications? The NMFEC already compresses raster data for remote Versatec printers and Graphics User Service Station television monitors. Is there more that can be done?

#### 8. Keeping Programs Up To Date

Distributed applications mean there are now two or more programs running on separate computers that must be kept up to date. Part of a distributed program is now in the hands of the user rather than the traditional case of a utility routine being completely within the control of the computer center. When the center updates such a distributed program, there must be a mechanism to detect obsolete programs at users' workstations and to update them automatically.

### 14.7. RESEARCH QUESTIONS

Many areas in distributed computing are still research questions. Some have been looked at, some are being looked at, but none are resolved.

The area of remote procedure calls needed to support distributed applications has been looked at by Sun Microsystems. Their proposed protocol, RPC, should be reviewed to see if it, or some equivalent, is the answer.

A unified file system is one that allows all computers access to any file, regardless of where the file resides. The file is accessible as though it were on the computer accessing the file. This is an attractive concept for distributed applications. Many approaches to this topic are already being pursued. The NMFEC is looking at a system-unifying Cray Timesharing System (CTSS) and the Common Filing System (CFS); Sun Microsystems has a Network File System (NFS); UNIX has its approach; and there are probably more. We need to see what is best for us.

Internetworking has already been raised as a potential problem. Users may be using Ethernet and DECnet or TCP/IP. Is there any graceful way to write distributed applications so that reasonable subsets of distributed functions would still be available depending on the various protocols in the chain between the user's workstation and the supercomputer?

The tools mentioned above for distributed applications programming generally do not exist. Furthermore, it is not well understood how to create them. Unfortunately, they will probably not be developed until people have acquired a lot of painful experience in distributed programming without the aid of tools.

## 14.8. SUMMARY

Distributed functions are computer applications that involve two or more closely integrated programs running simultaneously on separate computers. Distribution refers to the fact that what behaves as a single program is in fact several programs distributed over two or more computers. The most important case is exactly two computers, one of which is a supercomputer and the other of which is a workstation in the user's office. The separate computers must be connected by a hardware and software system that we call the network. A network can be geographically widespread.

Distributed applications clearly will make demands upon the network to move varying amounts of data between hosts at high speed. However, between the network level and the application level, there are several additional levels of software needed to support applications. Although these additional layers are not intrinsically a part of the network, they determine to a large extent how well distributed applications will work. The impact of distributed applications on the network itself is relatively clear; the impact on the higher layers of networking software raises considerably more questions than answers.

## 15. NATIONAL NETWORK REQUIREMENTS: A LOS ALAMOS PERSPECTIVE ON COMPUTER GRAPHICS AND DISTRIBUTED COMPUTING IN A SUPERCOMPUTER ENVIRONMENT

---

Robert L. Judd, Charles A. Slocomb, John F. Morrison,  
Dennis V. Brockway, and Bruce Eric Brown  
Los Alamos National Laboratory

### Abstract

*A primary goal of the Los Alamos Computing and Communications Division is to provide tools and capabilities that increase the productivity of our users. To achieve this goal, we have built a network that allows users to interactively access the central supercomputing systems using communications bandwidths that vary from 9600 bit/s to 40 Mbit/s. We have traditionally emphasized a timesharing supercomputing environment with high-speed graphics display capability. We are evolving toward a distributed supercomputing environment that uses high-performance workstations with more natural software interfaces to reduce the obstacles to man-machine communications. This environment will, to a large extent, integrate the software development and the graphics display activities. In addition, we are experimenting with ultra-high-speed graphic capabilities. These emerging technologies, together with our growing national user community, imply a hierarchy of bandwidth requirements for a network that is available to organizations from across the United States. These technologies and requirements are discussed in this paper.*

### 15.1. LOS ALAMOS COMPUTING ENVIRONMENT

The Los Alamos Central Computing Facility (CCF) consists of eight Cray computers, each running an interactive operating system, representing 22 CRAY-1 compute equivalents interconnected with a packet-switched network. Two-thirds of the supercomputing power is in a secure environment and one-third in an open environment. There are shared resources for graphics and printed output, file storage, and a terminal network. Computers in both the secure and open environments have access to the other resources in the CCF. Figure 1 is a representation of the CCF.

To satisfy different user requirements, we provide a variety of communications capabilities. These are shown in Table 1.

40 Mbit/s	Main CCF network connection
10 Mbit/s	Local area network (LAN) connections
330 kbit/s	High-speed serial connections
9600 bit/s	Serial connections

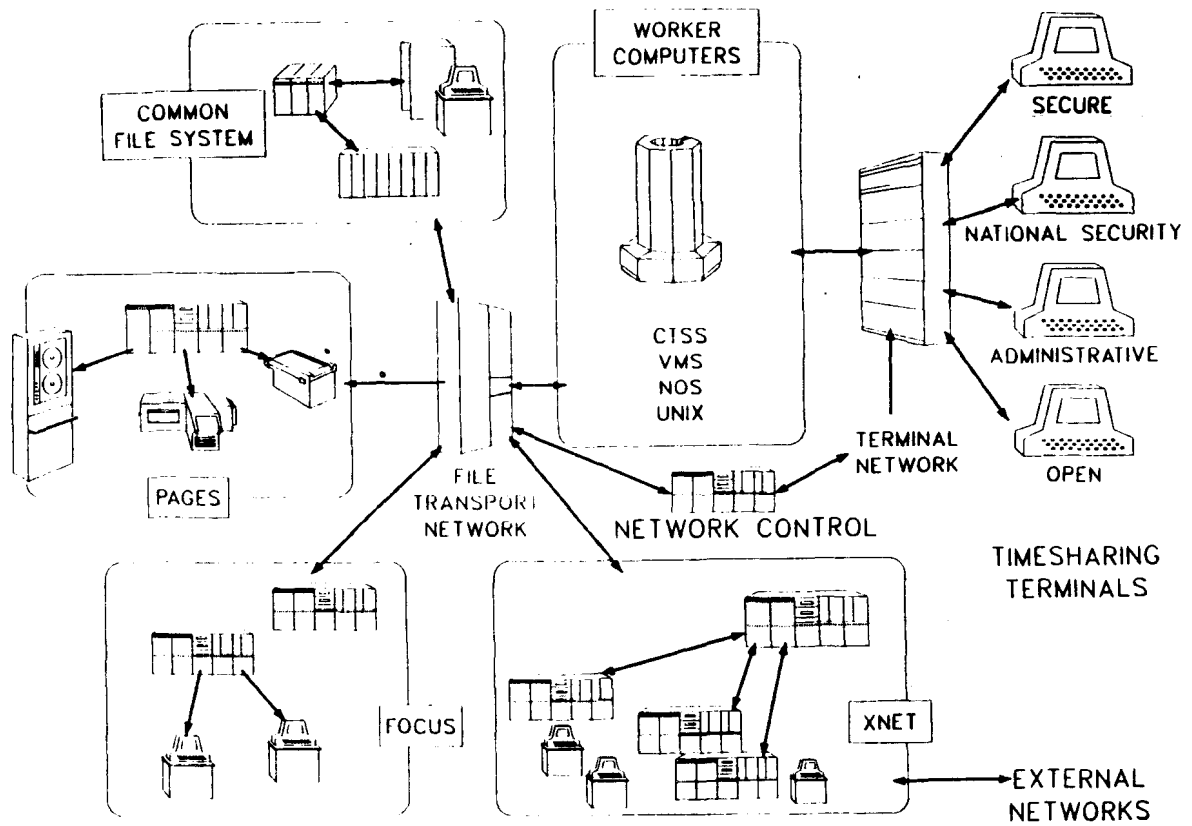


Figure 1. The Los Alamos Central Computing Facility.

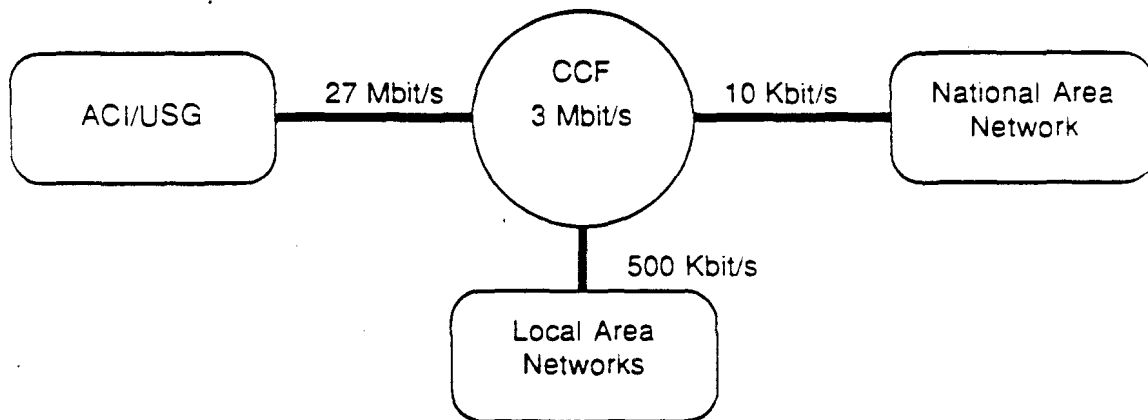


Figure 2. Current CCF and delivered data network speeds.

Most users of the supercomputers at Los Alamos communicate at 9600 bit/s (the effective data rate is 4800 bit/s for many terminals, but is 9600 bit/s for Tektronix graphics terminals). This line speed is the lowest bandwidth commonly available to users. It is used to access all the computing services including text editing, mail, program development and debugging, and some graphics needs. At this rate we can display a screen of data in about 2 s. Low-resolution raster images about 320 s to display.

The next level of communication speed is 330 kbit/s (with an effective data rate of 150 kbit/s). This service is provided to users housed adjacent to the CCF who have a need to view graphics data from supercomputer calculations. At this rate, a page of text is displayed in 0.13 s, point plots can be made at almost 4000 points per second, and vector plot speeds range from 2000 to 6000 vectors per second.

Local area networks are connected to the CCF through gateways and internet routers to 10-Mbit/s (Ethernet) lines and 5-Mbit/s broadband connections. The effective data rate is about 500 kbit/s. At this rate, a page of text can be displayed in 0.004 seconds, vectors can be drawn at 7000 to 30,000 vectors per second, and raster images can be displayed at about 12 seconds with 1024-bits by 768-bits by 8-bits resolution. We also have one user connected to a supercomputer at 40 Mbit/s. This application is described in more detail in a later section.

The current state of networks at Los Alamos is shown in Figure 2. The approximate speed of the links delivered by the CCF is indicated on the figure. Within the CCF we have file transfer and communications delivering data from 1 to 3 Mbit/s between worker computers.

Of particular concern to Los Alamos is security: both classified and unclassified information must be protected from unauthorized access. Any system used must be securable, and each subsystem added to the CCF must pass a rigorous security review.

## 15.2. PRODUCTIVITY

At Los Alamos we are convinced that people are more productive in an environment in which they can use the supercomputer systems interactively. Productivity of scientists and engineers who use the computer for creative processes can be increased by decreasing the response time of the computer system. In a study by IBM,<sup>1</sup> it was shown that productivity can be increased twofold by reducing the system response time from 2.0 to 0.3 s. In a later study, again by IBM,<sup>2</sup> the interactive productivity of engineers using a graphics program increased fourfold when the response time was reduced from 1.0 s to less than 0.3 s (see Figure. 3). After several years of study, James Brady of IBM found a correlation between improved transaction rate and productivity.<sup>3</sup> He also found a strong indication that an individual's error rate decreases as response time decreases. In his study he showed an expert engineer who operated on a "roll" during most of his work; that is, the time between successive steps of the process was minimized, and the attention of the individual was not distracted. So, as the ideas and solutions come to him, they suggest more ideas and solutions in rapid succession. When the computer response slows, this "roll" is interrupted and the productivity declines. Our goal is for our users to be on this "roll" so that they can work most productively.

We are approaching this productivity goal in two ways: we are moving toward a distributed supercomputing environment in which users have a workstation that is dedicated to their needs and therefore can provide a high level of responsiveness, and we are providing high-speed data communications to users who do extensive graphics. We are also experimenting with ultra-high-speed connections to a graphics system for animation. These two approaches are described below. We then discuss the implications of these technologies for national networking needs.




Figure 3. Productivity as a function of response time.

### 15.3. COMPUTER GRAPHICS

Computer graphics plays an important role in improving the productivity of the computer user. It allows the user to visualize the results of his computation instead of dealing with pages and pages of numbers. The time it takes to display the graphics generated from a calculation has a dramatic effect on the productivity of the user. By decreasing the time it takes to display a graphical image, we can free the user from the limitations of his tools, allowing him to concentrate on his science and not be distracted. All this is an attempt to allow the user to gain insight into his science. At Los Alamos we are trying to provide such an environment in several different ways.

The Advanced Computing Initiative/Ultra-High-Speed Graphics (ACI/USG) project at Los Alamos is an internally funded Institutional Supporting Research and Development project. The project's main purpose is to create an environment where computational physicists can increase their understanding of scientific problems and we can experiment with high-speed data communication options. The facility joins high-speed color graphics with an increased bandwidth connection to the supercomputer. All components have been based on commercially available technology in the first phase, but during the next phase we plan to develop components that push the limits of available technology in providing this capability.

The major components (see Figure 4) are a high-performance color graphic subsystem, real-time video disks, a high-performance minicomputer that will help process the graphics images, and a data communication connection to a supercomputer in the CCF. Included also are scientific workstations connected with Ethernet and personal computers for program development.



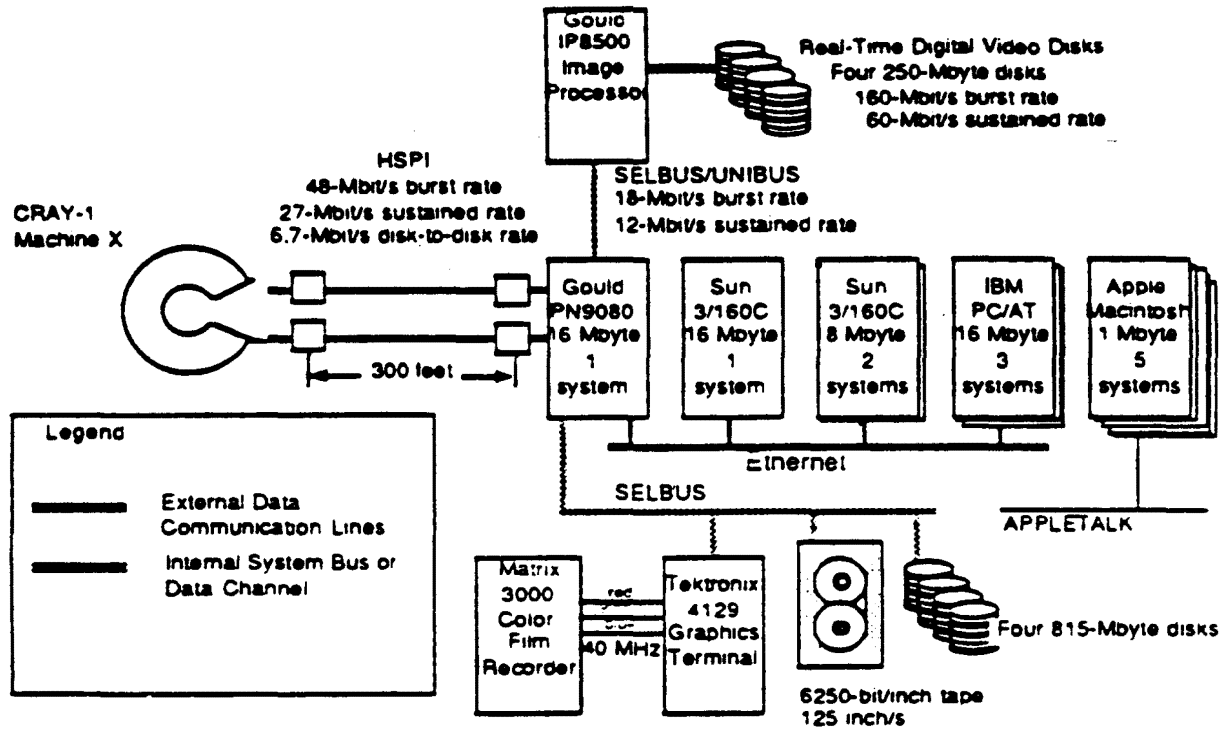


Figure 4. The ACI/USG facility.

Our initial physics application generates graphical images on the supercomputer and then moves these images to the Gould Image Support Processor (ISP) for processing. The images are then loaded on the real-time digital video disks where they can be displayed in an repeated animation sequence. The supercomputer calculation generates data in a 512- by 256- by 8-bit (1-Mbit) image that has a horizontal axis of symmetry. Once these images are moved to the ISP, they are processed to produce a 512- by 512- by 8-bit (2-Mbit) image by duplicating the data across the axis of symmetry and loaded on the digital video disks. The images are then played in a movie sequence that results in over 62.9 Mbit/s of user data throughput. The current 40-Mbit/s supercomputer connection delivers 27 Mbit/s of user throughput. This is a good match with the existing data being moved from the supercomputer. The current bottleneck is in the process of loading the images on the video disks that currently takes 0.9 s (2.33 Mbit/s) per image. We expect to eliminate this bottleneck by a new interface in early 1987. This work has provided us with a facility for experimentation to investigate the requirements for "real-time" interaction. The system would require a sustained data rate of 240 Mbit/s to display a 1024- by 1024- by 8-bit image at 30 times a second.

This system is being used for viewing animation sequences; therefore the IBM numbers are not applicable to this system. What is important is the bandwidth required to sustain the animation. The current 512- by 512- by 8-bit images require a data rate of 62.9 Mbit/s, while the 1024- by 1024- by 8-bit images would require a data rate of 240 Mbit/s. In the future we envision a user wanting to display 2048- by 2048- by 24-bit (100-Mbit) images at least 30 times a second. This requires a throughput of over 3 Gbit/s, or close to four times faster than the fastest external data channel on a Cray supercomputer. This illustrates that graphics will continue to stress the limits of data communications technology for the foreseeable future. The ACI/USG project will continue to explore ways to reduce the limitations of technology on the user. One of our key interest areas is to continue to develop high-bandwidth connections to supercomputers.

Another approach to increasing a user's productivity is to provide highly interactive graphical workstations in a distributed environment. The issue of whether a graphics workstation or a graphics terminal is the better tool is still being debated. Each has its strengths and weaknesses, but clearly both will require increased bandwidth in their connection to the remote supercomputer to reduce the response time of interactions with the supercomputer. We feel that a single communication path can serve the needs of both types of systems.

Technology has already provided graphics hardware that can outperform the same graphics generated in software on the fastest supercomputers. This trend will continue as technology provides higher performance graphical hardware for a reasonable price. Yet to exploit this increase in hardware performance from programs on a supercomputer, we will need new graphics software that allows the hardware to be exploited while reducing the amount of data to be transferred between these systems.

A national computer network must support and endorse a common set of graphic standards that would allow this distributed graphics environment to flourish and also provide the highest data throughput that is technologically feasible for a reasonable cost.

#### 15.4. DISTRIBUTED PROCESSING

Large-scale processors seem to be approaching the limit of their capabilities. Microcomputers, though, are still rapidly improving in both functionality and speed. Speeds of 4 to 10 million instructions per second (MIPS) are available on today's microcomputer systems. Their speed is doubling each year and this trend is expected to continue. Optimistic observers predict that 100-MIPS processors will be available in the next decade. These speeds are comparable to the instruction rates of today's fastest processors.

These microprocessors are being incorporated into workstations that combine fast computation rates with high-resolution monitors, large local memories, extensive disk storage capabilities, networking, local graphic hardware, and high-quality output devices. Perhaps more important than the enhanced hardware capabilities is the extensive software being developed in the mass market for these systems. This software includes programming environments that embody the most sophisticated screen editors, debuggers, software engineering tools for design and documentation of application codes in Fortran and other languages, and software to support high-speed graphics capabilities.

Although, on the one hand, these workstations are becoming powerful computational devices in their own right, they will never match the scientific computational capability of supercomputers. On the other hand, supercomputers have never included the software development environment needed to make code development people most effective. What is required is a computing environment that exploits the best of both of these systems in a distributed supercomputing environment.

This distributed environment will use the workstation as a transparent front end to the centrally located supercomputer systems. The user will see the window-oriented monitor of the workstation as the primary interface to the supercomputer. The interactive portion of the tools and application codes will run on the workstation, whereas the more computationally intensive part of the code will run on the supercomputer. Programs will be developed with remote procedure calls (RPCs) over the network to distribute the processing. The user, neither knowing nor caring where the tools are running, will appear to be on the workstation and will benefit from the response that a dedicated workstation will provide. The details of the operating system will become less important to the users with these windowing tools. These features will allow a larger group of people to benefit because it will no longer be necessary to be an expert on a variety of operating systems to be productive.

Together with the development of workstation hardware and software, there are continuing advances in networking technology that allow centralized services and high-speed communications between concentrations of users at various sites. This capability creates new possibilities for workstation users. In addition to providing the distributed supercomputing environment described above, possibilities exist for widespread access to other Laboratory services, such as mail, library services, and engineering databases, from the user's workstation.

For this scenario to occur we need standards in several areas. At the lowest level we need standard network protocols. The TCP/IP standard fills the lower level needs now. In addition, we need a standard network file system and window system. A network file system allows the interoperability of a program or project across possibly heterogeneous operating systems. Programs then have transparent access to remote or foreign file systems. Currently there are two possible candidates for this function. They are the Sun Microsystem's Network File System (NFS) and AT&T's Remote File System (RFS).<sup>4</sup>

One window standard is the network extensible window system (NeWS) recently announced by Sun Microsystems. NeWS allows the workstation screen to have fully integrated graphics and text driven by device-independent remote programs. This feature allows the easy display and manipulation of complex images on the workstation. The workstation is no longer limited to a textual standard display method but is able to display pictures wherever they come from, with the sending device knowing nothing about the workstation. Another possible standard is the Massachusetts Institute of Technology X-Window system. Both will allow an intimate integration of supercomputers with the workstations in offices. We believe that data throughputs of at least 100 kbit/s are required to support a reasonable level of interactivity in a distributed computing environment.

### 15.5. REQUIREMENTS OF A NATIONAL NETWORK

We believe that there is a demonstrable need for the Los Alamos staff to access the resources available through a national network, and there is also a need for people around the nation to access the resources at Los Alamos. Figure 5 shows the networks currently connected to the CCF with the required bandwidths and effective data rates of these connections.

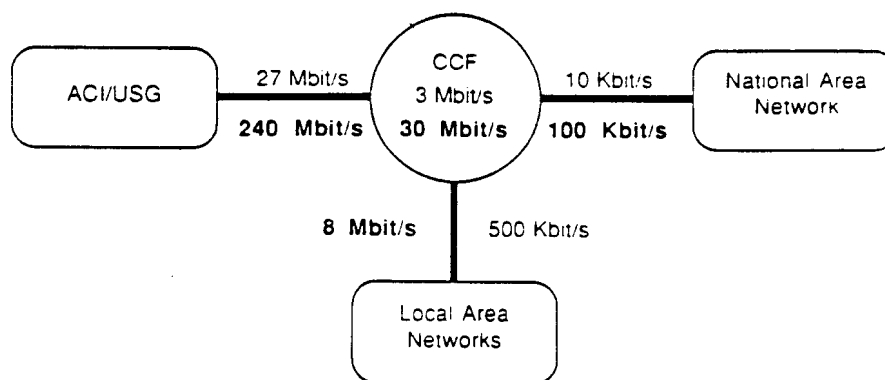


Figure 5. Los Alamos system future connection speeds.

## 15.6. SUMMARY

At Los Alamos, we are evolving toward a distributed supercomputing environment that will provide users with high-bandwidth connections to the CCF. We currently provide local users with a variety of communications capabilities to meet different computing requirements, and we are experimenting with ultra-high-speed communications for graphics. In addition to the local users, we have remote users all over the United States who have similar computing requirements. Our basic need for access to a national network is to satisfy the computing requirements of remote users, to enable scientists within the Laboratory to collaborate easily with scientists around the nation and around the world, to allow Laboratory staff and other scientists to access national databases, and to allow access to experimental computer architectures. Network characteristics to meet these requirements are given below.

- The speed of the national network is of crucial importance. The Los Alamos CCF now has a national user base. We need high bandwidths so that these remote users will also have a productive environment. For a productive distributed supercomputing capability we require burst data throughputs of at least 500 kbit/s. For a productive graphics capability, we will need burst data rates of at least 1-Mbit/s throughput for each user. We believe that the initial national network connection to Los Alamos should be at least the T1 bandwidth<sup>5</sup> and over the next 5 to 10 years should increase an order of magnitude. Our current ARPANET connection provides the Laboratory with 56 kbit/s. This bandwidth is adequate for some functions such as mail, but it is not adequate for interactive computing use.
- The scientific community with whom Laboratory scientists collaborate is both national and international in scope. Therefore, it is important that the network be interconnected with other prominent national and international networks. It should be accessible from our facility with a single interface using standard protocols.
- The network should be accessible from major cities around the nation so that travelers can access it from wherever they are.
- Electronic mail will include digitized and near-typeset-quality images that are stored at 300 dots per inch. An 8.5- by 11-inch sheet digitized at 300 dots per inch produces about 1 Mbyte (8 Mbits) of data. Similarly, voice digitized at 64 kHz, sampled with 8 bits for 30 s, results in 1.966 Mbyte of data. While these numbers represent large amounts of data, the speed that they need to be transmitted is not as critical as the graphic interactivity described above. We also see the need for transmission of live NTSC video. Some teleconferencing systems today run with a data rate of 1.5 Mbit/s. True digital TV will require 45 Mbit/s.
- Half of the communications with the CCF at Los Alamos are secured lines. Any national network will have to handle security such that classified work may be done remotely and unauthorized access to any information must be prevented.

## 15.7. REFERENCES

1. W. J. Doherty and R. P. Kelishy, "Managing VM/CMS Systems for User Effectiveness," *IBM Systems Journal*, **18**, 1, pp. 143-163 (1979).
2. "The Economic Value of Rapid Response Time," International Business Machines Corporation (1982).
3. J. T. Brady, "A Theory of Productivity in the Creative Process," *Proceedings of the IEEE 1st International Conference on Computer Workstations*, pp. 70-79 (November 1985).
4. M. J. Hatch, M. Katz, and J. Rees, "AT&T's RFS and Sun's NFS," *Unix/World*, pp. 40-51 (December 1985).
5. J. E. McNamara, *Technical Aspects of Data Communications*, Second Edition, Digital Press, Bedford, Massachusetts (1982).



## 16. NETWORK REQUIREMENTS FOR SCIENTIFIC RESEARCH

---

**Barry M. Leiner**

Research Institute for Advanced Computer Science\*  
NASA Ames Research Center

### **Abstract**

*Computer networks are critical to scientific research. The recognition of that fact has prompted several agencies to fund networks for their researchers. This workshop is aimed at investigating the cooperation between these agencies to in order to provide these functions to the broad scientific community in a cost-effective manner. This paper attempts to outline the requirements for such a national research internetwork. It first addresses the functions a user requires of a network and then addresses near-term requirements and future goals for such a network.*

### **16.1. INTRODUCTION**

Computer networks are critical to scientific research. They are currently being used by portions of the scientific community to support access to remote resources (such as supercomputers and data at collaborator's sites) and collaborative work through such facilities as electronic mail and shared databases. There is considerable movement in the direction of providing these capabilities to the broad scientific community in a unified manner, as evidenced by this workshop. In the future, these capabilities will even be required in space, as the space station becomes a reality as a scientific research resource.

The purpose of this paper is to identify the range of requirements for networks that are to support scientific research. These requirements include the basic connectivity provided by the links and switches of the network through the basic network functions to the necessary user services to allow effective use of the interconnected network. The paper has four sections. The first section discusses the functions a user requires of a network. The second section discusses the requirements for the underlying link and node infrastructure, while the third proposes a set of specifications to achieve the functions on an end-to-end basis. The fourth section discusses a number of network-oriented user services that are needed in addition to the network itself. In each section, the discussion is broken into two categories. The first addresses near-term requirements: those capabilities and functions that are needed today and for which technology is available to perform the function. The second category addresses long-term goals: those capabilities for which additional research is needed.

### **16.2. NETWORK FUNCTIONS**

This section addresses the functions and capabilities that networks and particularly internetworks should be expected to support in the near-term future.

---

\*Work reported herein was supported by Cooperative Agreement NCC 2-387 from the National Aeronautics and Space Administration to the Universities Space Research Association. This paper was written by the IAB Task Force on Scientific Computing and edited by Dr. Leiner.

## 16.2.1. Near-Term Requirements

There are many functions that are currently available to subsets of the user community. These functions should be made available to the broad scientific community.

### 16.2.1.1. User/Resource Connectivity

Undoubtedly the first order of business in networking is to provide interconnectivity of users and the resources they need. The goal in the near term for internetworking should be to extend the connectivity as widely as possible, i.e., to provide ubiquitous connectivity among users and between users and resources. Note that the existence of a network path between sites does not necessarily imply interoperability between communities and or resources using noncompatible protocol suites. However, a minimal set of functions should be provided across the entire user community, independent of the protocol suite being used. These typically include electronic mail at a minimum; file transfer and remote login capabilities must also be provided.

### 16.2.1.2. Home Usage

One condition that could enhance current scientific computing would be to extend to the home the same level of network support that the scientist has available in his office environment. As network access becomes increasingly widespread, the extension to the home will allow the user to continue his computing at home without dramatic changes in his work habits, based on limited access.

### 16.2.1.3. Charging

The scientific user should not have to worry about the costs of data communications any more than he worries about voice communications (his office telephone), so that data communications become an integral and low-cost part of our national infrastructure. This implies that charges for network services must *not* be volume sensitive and must *not* be charged back to the individual. Either of these conditions forces the user to consider network resources as scarce and therefore as requiring his individual attention to conserve them. Such attention to extraneous details not only detracts from the research, but fundamentally impacts the use and benefit that networking is intended to supply. This does not require that networking usage is free. It should be either be low enough cost that the individual does not have to be accountable for "normal" usage or managed in such a manner that the individual does not have to be concerned with it on a daily basis.

### 16.2.1.4. Applications

Most applications in the near term that must be supported in an internetwork environment are essentially extensions of current ones.

- **Electronic Mail**--Electronic mail will increase in value as the extended interconnectivity provided by internetworking makes users more reachable.
- **Multimedia Mail**--An enhancement to text-based mail will include capabilities such as figures, diagrams, graphs, and digitized voice.
- **Multimedia Conferencing**--Network conferencing is communication among multiple people simultaneously. Conferencing may or may not be done in real time; that is, all participants may not be required to be online at the same time. The multimedia supported may include text, voice, video, graphics, and possibly other capabilities.



- **File Transfer**--The ability to transfer data files.
- **Bulk Transfer**--The ability to stream large quantities of data.
- **Interactive Remote Login**--The ability to perform remote terminal connections to hosts.
- **Remote Job Entry**--The ability to submit batch jobs for processing to remote hosts and receive output.

Applications that need support in the near term but are *not* extensions of currently supported applications include the following.

- **Remote Instrument Control**--This normally presumes to have a "human in the control loop." This condition relaxes the requirements on the (inter)network somewhat as to response times and reliability. Timing would be presumed to be commensurate with human reactions and reliability would not be as stringent as that required for completely automatic control.
- **Remote Data Acquisition**--This supports the collection of experimental data where the experiment is remotely located from the collection center. This requirement can only be satisfied when the bandwidth, reliability, and predictability of network response are sufficient. This cannot be supported in the general sense because of the enormous bandwidth, very high reliability, and/or guaranteed short response time required for many experiments.

These last two requirements are especially crucial when one considers remote experimentation such as will be performed on the space station.

#### 16.2.1.5. Capabilities

The above applications could be best supported on a network with infinite bandwidth, zero delay, and perfect reliability. Unfortunately, even currently feasible approximations to these levels of capabilities can be very expensive. Therefore, it can be expected that compromises will be made for each capability and between them, with different balances struck between different networks. Because of this, the user must be given an opportunity to declare which capability or capabilities is/are of most interest--most likely through a "type-of-service" required declaration. Some examples of possible tradeoffs are the following.

- **File Transport**--This tradeoff normally requires primarily high reliability and high bandwidth secondarily. Delay is not as important.
- **Bulk Transport**--Some applications such as digitized video might require high bandwidth as the most important capability. Depending on the application, delay would be second and reliability would be of lesser importance. Image transfers of scientific data sometimes will invert the latter two requirements.
- **Interactive Traffic**--This normally requires low delay as a primary consideration. Reliability may be secondary depending on the application. Bandwidth would usually be of least importance.

#### 16.2.1.6. Standards

The use of standards in networking is directed toward interoperability and availability of commercial equipment. However, as stated earlier, full interoperability across the entire scientific community is probably not a reasonable goal for internetworking in the near term because of the protocol mix now present. That is not to say, though, that the use of standards should not be pursued on the path to full user interoperability. Standards, in the context of near-term goal support, include the following.

- **Media Exchange Standards**--These standards would allow the interchange of equations, graphics, images, and databases, as well as text.
- **Commercially Available Standards**--Plug-compatible, commercially available standards will allow a degree of interoperability prior to the widespread availability of the International Standards Organization (ISO) standard protocols.

### 16.2.2. Long-Term Goals

In the future, the internetwork should be transparent communications between users and resources, and provide the additional network services required to make use of those communications. A user should be able to access whatever resources are available just as if the resources are in the user's office. The same high level of service should exist independent of which network one happens to be on. In fact, one should not even be able to tell that the network is there!

It is also important that people be able to work effectively while at home or when traveling. Wherever one may happen to be, it should be possible to "plug into" the internetwork and read mail, access files, control remote instruments, and have the same kind of environment one is used to at the office.

Services to locate required facilities and take advantage of them must also be available on the network. These range from the basic White and Yellow pages, providing network locations (addresses) for users and capabilities, to distributed databases and computing facilities. Eventually, this conglomeration of computers, workstations, networks, and other computing resources will become one gigantic distributed "world computer" with a very large number of processing nodes all over the world.

## 16.3. NETWORK CONNECTIVITY

By network connectivity we mean the ability to move packets from one point to another.\* Note that this need not mean functional interoperability, since the endpoints may be using incompatible protocols. Thus, in this section, we will be addressing the use of shared links and interconnected networks to provide a possible path. In the next section, the exploitation of these paths to achieve functional connectivity will be addressed.

In this section, we discuss the need for providing these network paths to a wide set of users and resources, and the characteristics of those paths. As in other sections, this discussion is broken into two major categories. The first category contains those goals that we believe to be achievable with currently available technology and implementations. The second category is those areas in which further research is required.

### 16.3.1. Near-Term Objectives

Currently, there are a large number of networks serving the scientific community including ARPANET, MFEnet, SPAN, NASnet, and the NSFnet backbone. Although there is some loose correlation between the networks and the disciplines they serve, these networks are organized more based on federal funding. Furthermore, while there is significant interconnectivity between a number of the networks, there is considerable room for more sharing of these resources.

---

\*Note that an implicit assumption in this paper is that packet switched networks are the preferred technology for providing a scientific computer network. This is due to the ability of such networks to share the available link resources to provide interconnection between numerous sites and their ability to effectively handle the "bursty" computer communication requirement.

In the near term, therefore, there are two major requirement areas: providing for connectivity based on discipline and user community, and providing for the effective use of adequate networking resources.

#### 16.3.1.1. Discipline Connectivity

Scientists in a particular community/discipline need to have access to many common resources as well as communicate with each other. For example, the quantum physics research community obtains funding from a number of federal sources but carries out its research within the context of a scientific discourse. Furthermore, this discourse often overlaps several disciplines. Because networks are generally oriented based on the source of funding, this required connectivity has in the past been inhibited. NSFnet is a major step towards satisfying this requirement, because of its underlying philosophy of acting as an interconnecting network between supercomputer centers and between state, regional, and campus networks. This move towards a set of networks that are interconnected, at least at the packet transport level, must be continued so that a scientist can obtain connectivity between his/her local computing equipment and the computing and other resources that are needed independently of the source of funds.\*

#### 16.3.1.2. Communication Resource Sharing

The scientific community is always going to suffer from a lack of adequate communication bandwidth and connections. There are requirements (e.g., graphic animation from supercomputers) that stretch the capabilities of even the most advanced long-haul networks. In addition, as more and more scientists require connection into networks, the ability to provide those connections on a network-centric basis will become more and more difficult.

However, the communication links (e.g., leased lines and satellite channels) providing the underlying topology of the various networks span in aggregate a very broad range of the scientific community sites. If, therefore, the networks could share these links in an effective manner, the following two objectives could be achieved.

- (1) The need to add links just to support a particular network topology change would be decreased.
- (2) New user sites could be connected more readily.

Existing technology (namely the DARPA-developed gateway system based on the Internet Protocol, IP) provides an effective method for accomplishing this sharing. By using IP gateways to connect the various networks and by arranging for suitable cost sharing, the underlying connectivity would be greatly expanded and both of the above objectives achieved.

#### 16.3.1.3. Expansion of Physical Structure

Unfortunately, the mere interconnectivity of the various networks does not increase the bandwidth available. While it may allow for more effective use of that available bandwidth, a sufficient number of links with adequate bandwidth must be provided to avoid network congestion. This problem has already occurred in the ARPANET, where the expansion of the use of the network without a concurrent expansion in the trunking and topology has resulted in congestion and consequent degradation in performance.

---

\*Obviously, actual use of those resources will depend on obtaining access permission from the appropriate controlling organization. For example, use of a supercomputer will require permission and some allocation of computing resources. The lack of network access should not, however, be the limiting factor for resource utilization.

Thus, it is necessary to augment the current physical structure (links and switches) both by increasing the bandwidth of the current configuration and by adding additional links and switches where appropriate.

#### 16.3.1.4. Network Engineering

One of the major deficiencies in the current system of networks is the lack of overall engineering. While each of the various networks generally is well supported, there is woefully little engineering of the overall system. As the networks are interconnected into a larger system, this need will become more severe. Examples of the areas where engineering is needed are as follows.

- **Topology Engineering**--This area involves deciding where links and switches should be installed or upgraded. If the interconnection of the networks is achieved, this will often involve a decision as to which networks need to be upgraded as well as deciding where in the network those upgrades should take place.
- **Connection Engineering**--When a user site desires to be connected, which node of which network is the best for that site must be determined, and such issues as existing node locations, available bandwidth, and expected traffic patterns to/from that site must be considered.
- **Operations and Maintenance**--This area involves monitoring the operation of the overall system and identifying corrective actions when failures occur.

#### 16.3.1.5. Support of Different Types of Service

Several different end user applications are currently in place, and these put different demands on the underlying structure. For example, interactive remote login requires low delay, while file transfer requires high bandwidth. It is important in the installation of additional links and switches that care be given to providing a mix of link characteristics. For example, high bandwidth satellite channels may be appropriate to support broadcast applications or graphics, while low delay will be required to support interactive applications.

#### 16.3.2. Future Goals

Significant expansion of the underlying transport mechanisms will be required to support future scientific networking. These expansions will be both in size and performance.

##### 16.3.2.1. Bandwidth

Bandwidth requirements are being driven higher by advances in computer technology as well as the proliferation of that technology. As high-performance graphics workstations work cooperatively with supercomputers, and as real-time remote robotics and experimental control become a reality, the bandwidth requirements will continue to grow. In addition, as the number of sites on the networks increase, so will the aggregate bandwidth requirement. However, at the same time, the underlying bandwidth capabilities are also increasing. Satellite bandwidths of tens of megabits are available, and fiber optics technologies are providing extremely high bandwidths (in the range of gigabits). It is therefore essential that the underlying connectivity take advantage of these advances in communications to increase the available end-to-end bandwidth.

### 16.3.2.2. Expressway Routing

As higher levels of internet connectivity occur, there will be a new set of problems related to lowest hop count and lowest delay routing metrics. The assumed internet connectivity can easily present situations where the highest speed, lowest delay route between two nodes on the same net is via a route on another network. Consider two sites one either end of the country, but both on the same multipoint internet, where their network also is gatewayed to some other network with high-speed transcontinental links. The routing algorithms must be able to handle these situations gracefully, and they become of increased importance in handling global type-of-service routing.

## 16.4. NETWORK SPECIFICATIONS

To achieve the end-to-end user functions discussed in Section 16.2, it is not adequate to simply provide the underlying connectivity described in the previous section. The network must provide a certain set of capabilities on an end-to-end basis. In this section, we discuss the specifications on the network that are required.

### 16.4.1. Near-Term Specifications

In the near term, the requirements on the networks are twofold. The first is to provide those functions that will permit full interoperability. Second, the internetwork must address the additional requirements that arise in the connection of networks, users, and resources.

#### 16.4.1.1. Interoperability

A first-order requirement for scientific computer networks (and computer networks in general) is that they be interoperable with each other, as discussed in the above section on connectivity. A first step to accomplish this is to use IP. The use of IP will allow individual networks built by differing agencies to combine resources and minimize cost by avoiding the needless duplication of network resources and their management. However, use of IP does not provide end-to-end interoperability. There must also be compatibility of higher level functions and protocols. At a minimum, while commonly agreed upon standards (such as the ISO developments) are proceeding, methods for interoperability between different protocol suites must be developed. This would provide interoperability of certain functions, such as file transfer, electronic mail, and remote login. The emphasis, however, should be on developing agreement within the scientific community on use of a standard set of protocols.

#### 16.4.1.2. Access Control

The design of the network should include adequate methods for controlling access to the network by unauthorized personnel. This especially includes access to network capabilities that are reachable via the commercial phone network and public data nets. For example, terminal servers that allow users to dial up via commercial phone lines should have adequate authentication mechanisms in place to prevent access by unauthorized individuals. However, it should be noted that most hosts that are reachable via such networks are also reachable via other "non-network" means, such as directly dialing over commercial phone lines. The purpose of network access control is not to insure isolation of hosts from unauthorized users, and hosts should not expect the network itself to protect them from "hackers."

### 16.4.1.3. Privacy

The network should provide protection of data that traverses it in a way that is commensurate with the sensitivity of that data. It is judged that the scientific requirements for privacy of data traveling on networks do not warrant a large expenditure of resources in this area. However, nothing in the network design should preclude the use of link level or end-to-end encryption, or other such methods that can be added at a later time. An example of this kind of capability would be use of KG-84A link encryptors on MILNET or the Fig Leaf DES-based end-to-end encryption box developed by DARPA.

### 16.4.1.4. Accounting

The network should provide adequate accounting procedures to track the consumption of network resources. Accounting of network resources is also important for the management of the network, and particularly the management of interconnections with other networks. Proper use of the accounting database should allow network management personnel to determine the "flows" of data on the network and the identification of bottlenecks in network resources. This capability also has secondary value in tracking down intrusions of the network, and to provide an audit trail if malicious abuse should occur. In addition, accounting of higher level network services (such as terminal serving) should be kept track of for the same reasons.

### 16.4.1.5. Type of Service Routing

Type of service routing is necessary since not all elements of network activity require the same resources, and the opportunities for minimizing use of costly network resources are large. For example, interactive traffic such as remote login requires low delay so the network will not be a bottleneck to the user attempting to do work. Yet the bandwidth of interactive traffic can be quite small compared to the requirements for file transfer and mail service that are not response-time critical. Without type-of-service routing, network resources must be sized according to the largest user and have characteristics that are pleasing to the most finicky user. This has major cost implications for the network design, as high-delay links, such as satellite links, cannot be used for interactive traffic despite the significant cost savings they represent over terrestrial links. With type of service routing in place in the network gateways, and proper software in the hosts to make use of such capabilities, overall network performance can be enhanced, and sizable cost savings realized. Since the IP protocol already has provisions for such routing, such changes to existing implementations do not require a major change in the underlying protocol implementations.

### 16.4.1.6. Administration of Address Space

Local administration of network address space is essential to provide for prompt addition of hosts to the network, and to minimize the load on backbone network administrators. Further, a distributed name to address translation service also has similar advantages. The DARPA Name Domain system currently in use on the internet is a suitable implementation of such a name-to-address translation system.

### 16.4.1.7. Remote Procedure Call Libraries

In order to provide a standard library interface so that distributed network utilities can easily communicate with each other in a standard way, a standard Remote Procedure Call (RPC) library must be deployed. The computer industry has lead the research community in developing RPC implementations, and current implementations tend to be compatible within the same type of operating

system, but not across operating systems. Nonetheless, a portable RPC implementation that can be standardized can provide a substantial boost in present capability to write operating system independent network utilities. If a new RPC mechanism is to be designed from scratch, then it must have enough capabilities to lure implementors away from current standards. Otherwise, modification of an existing standard that is close to the mark in capabilities seems to be in order, with the cooperation of vendors in the field to assure implementations will exist for all major operating systems in use on the network.

#### **16.4.1.8. Remote Job Entry (RJE)**

The capabilities of standard network RJE implementations are inadequate and are implemented prolifically among major operating systems. While the notion of RJE evokes memories of dated technologies such as punch cards, the concept is still valid, and is favored as a means of interaction with supercomputers by science users. All major supercomputer manufacturers support RJE access in their operating systems, but many do not generalize well into the internet domain. That is, a RJE standard that is designed for 2400-baud modem access from a card reader may not be easily modifiable for use on the internet. Nonetheless, the capability for a network user to submit a job from a host and have its output delivered on a printer attached to a different host would be welcomed by most science users. Further, having this capability interoperate with existing RJE packages would add a large amount of flexibility to the whole system.

#### **16.4.1.9. Multiple Virtual Connections**

The capability to have multiple network connections open from a user's workstation to remote network hosts is an invaluable tool that greatly increases user productivity. The network design should not place limits (procedural or otherwise) on this capability.

#### **16.4.1.10. Network Operation and Management Tools**

The present state of internet technology requires the use of personnel who are, in the vernacular of the trade, called network "wizards," for the proper operation and management of networks. These people are a scarce resource to begin with, and squandering them on day-to-day operational issues detracts from progress in the more developmental areas of networking. The cause of this problem is that a good part of the knowledge for operating and managing a network has never been written down in any sort of concise fashion, and the reason for that is because networks of this type in the past were primarily used as a research tool, not as an operational resource. While the usage of these networks has changed, the technology has not adjusted to the new reality that a wizard may not be nearby when a problem arises. To insure that the network can flexibly expand in the future, new tools must be developed that allow non-wizards to monitor network performance, determine trouble spots, and implement repairs or "work-arounds."

#### **16.4.2. Future Goals**

The networks of the future must be able to support transparent access to distributed resources of a variety of different kinds. These resources will include supercomputer facilities, remote observing facilities, distributed archives and databases, and other network services. Access to these resources is to be made widely available to scientists, other researchers, and support personnel located at remote sites over a variety of internetworked connections. Different modes of access must be supported that are consonant with the sorts of resources that are being accessed, the data bandwidths required, and the type of interaction demanded by the application.

Network protocol enhancements will be required to support this expansion in functionality; mere increases in bandwidth are not sufficient. The number of end nodes to be connected is in the hundreds of thousands, driven by increasing use of microprocessors and workstations throughout the community. Fundamentally different sorts of services from those now offered are anticipated, and dynamic bandwidth selection and allocation will be required to support the different access modes. Large-scale internet connections among several agency size internets will require new approaches to routing and naming paradigms. All of this must be planned so as to facilitate transition to the ISO/OSI (Open Systems Interconnect) standards as these mature and robust implementations are placed in service and tuned for performance.

Several specific areas are identified below as being of critical importance in support of future network requirements. These are listed in no particular order.

#### **16.4.2.1. Standards and Interface Abstractions**

As more and different services are made available on these various networks, it will become increasingly important to identify interface standards and suitable application abstractions to support remote resource access. These abstractions may be applicable at several levels in the protocol hierarchy and can serve to enhance both applications functionality and portability. Examples are transport or connection layer abstractions that support applications independence from lower level network realizations or interface abstractions that provide a data description language that can handle a full range of abstract data type definitions. Applications or connection level abstractions can provide a means of bridging across different protocol suites as well as helping with protocol transition.

#### **16.4.2.2. OSI Transition and Enhancements**

Further evolution of the OSI network protocols and realization of large-scale networks so that some of the real protocol and tuning issues can be dealt with must be anticipated. It is only when such networks have been created that these issues can be approached and resolved. Type-of-service and Expressway routing and related routing issues must be resolved before a real transition can be contemplated. Using the interface abstraction approach just described will allow definition now of applications that can transition as the lower layer networks are implemented. Applications gateways and relay functions will be a part of this transition strategy, along with dual mode gateways and protocol translation layers.

#### **16.4.2.3. Processor Count Expansion**

Increases in the numbers of nodes and host sites and the expected growth in use of microcomputers, super-micro workstations, and other modest cost but high-power computing solutions will drive the development of different network and interconnect strategies as well as the infrastructure for managing this increased name space. Hierarchical name management (as in domain-based naming) and suitable transport layer realizations will be required to build networks that are robust and functional in the face of the anticipated expansions.

#### **16.4.2.4. Dynamic Binding of Names to Addresses**

Increased processor counts and increased usage of portable units, mobile units, and lap-top micros will make dynamic management of the name/address space a must. Units must have fixed designations that can be rebound to physical addresses as required or expedient.



## 16.5. USER SERVICES

The user services of the network are a key aspect of making the network directly useful to the scientist. Without the right user services, network users separate into artificial subclasses based on their degree of sophistication in acquiring skill in the use of the network. Flexible information dissemination equalizes the effectiveness of the network for different kinds of users.

### 16.5.1. Near-Term Requirements

In the near term, the focus is on providing the services that allow users to take advantage of the functions that the interconnected network provides.

#### 16.5.1.1. Directory-Services

Much of the information necessary in the use of the network is for directory purposes. The user needs to access resources available on the network, and needs to obtain a name or address.

- **White Pages**--The network needs to provide mechanisms for looking up names and addresses of people and hosts on the network. Flexible searches should be possible on multiple aspects of the directory listing. Some of these services are normally transparent to the user/host name to address translation, for example.
- **Yellow Pages**--Other kinds of information lookup are based on cataloging and classification of information about resources on the networks.

#### 16.5.1.2. Information Sharing Services

- **Bulletin Boards**--The service of the electronic bulletin board is the one-to-many analog of the one-to-one service of electronic mail. A bulletin board provides a forum for discussion and interchange of information. Accessibility is network-wide depending on the definition of the particular bulletin board. Currently the SMTP and UUCP protocols are used in the transport of postings for many bulletin boards, but any similar electronic mail transport can be substituted without affecting the underlying concept. An effectively open-ended recipient list is specified as the recipient of a message, which then constitutes a bulletin board posting. A convention exists as to what transport protocols are utilized for a particular set of bulletin boards. The user agent used to access the bulletin board may vary from host to host. Some number of host resources on the network provide the service of progressively expanding the symbolic mail address of the bulletin board into its constituent parts, as well as relaying postings as a service to the network. Associated with this service is the maintenance of the lists used in distributing the postings. This maintenance includes responding to requests from bulletin board readers and host bulletin board managers, as well as drawing the appropriate conclusions from recurring automatically generated error messages or error messages in response to distribution attempts.
- **Community Archiving**--Much information can be shared over the network. At some point, each particular information item reaches the stage where it is no longer appropriately kept online and accessible. When moving a file of information to offline storage, a network can provide its hosts a considerable economy if information of interest to several of them needs only be stored offline once. Procedures then exist for querying and retrieving from the set of offline stored files.

- **Shared/Distributed File System**--It should be possible for a user on the network to look at a broadly defined collection of information on the network as one useful whole. To this end, standards for accessing files remotely are necessary. These standards should include means for random access to remote files, similar to the generally employed on a single computer system.
- **Distributed Databases and Archives**--As more scientific disciplines computerize their data archives and catalogs, mechanisms will have to be provided to support distributed access to these resources. Fundamentally new kinds of collaborative research will become possible when such resources and access mechanisms are widely available.

### 16.5.1.3. Resource Sharing Services

In sharing the resources or services available on the network, certain ancillary services are needed depending on the resource.

- **Access Control**--Identification and authorization is needed for individuals, hosts, or subnetworks permitted to make use of a resource available via the network. There should be consistency of procedure for obtaining and utilizing permission for use of shared resources. The identification scheme used for access to the network should be available for use by resources as well. In some cases, this will serve as sufficient access control, and in other cases it will be a useful adjunct to resource-specific controls. The information on the current network location of the user should be available along with information on user identification to permit added flexibility for resources. For example, it should be possible to verify that an access attempt is coming from within a state. A state agency might then grant public access to its services only for users within the state. Attributes of individuals should be codifiable within the access control database, for example membership in a given professional society.
- **Privacy**--Users of a resource have a right to expect that they have control over the release of the information they generate. Resources should allow classifying information according to degree of access, i.e., none, access to read, access according to criteria specified in the data itself, and ability to change or add information. The full range of identification information described under access control should be available to the user when specifying access. Access could be granted to all fellow members of a professional society, for example.
- **Accounting**--To permit auditing of usage, accounting information should be provided for those resources for which it is deemed necessary. This would include identity of the user of the resource and the corresponding volume of resource components.

### 16.5.1.4. Legalities of Interagency Research Internet

To make the multiply sponsored internetwork feasible, the federal budget will have to recognize that some usage outside a particular budget category may occur. This will permit the cross-utilization of agency-funded resources. For example, NSFnet researchers would be able to access supercomputers over NASnet. In return for this, the total cost to the government will be significantly reduced because of the benefits of sharing network and other resources, rather than duplicating them.

### 16.5.1.5. Standards

In order for the networking needs of scientific computing to be met, new standards are going to evolve. It is important that they be tested under actual use conditions and that feedback be used to refine them. Since the standards for scientific communication and networking are to be experimented with, they are more dynamic than those in other electronic communication fields. It is critical that the resources of the network be expended to promulgate experimental standards and maximize the range of the community utilizing them. To this end, the sharing of results of the testing is important.

### **16.5.1.6. User-Oriented Documentation**

The functionality of the network should be available widely without the costly need to refer requests to experts for formulation. A basic information facility in the network should therefore be developed. The network should be self-documenting via online help files, interactive tutorials, and good design. In addition, concise, well-indexed, and complete printed documentation should be available.

### **16.5.2. Future Goals**

The goal for the future should be to provide advanced user services that allow full advantage to be taken of the interconnection of users, computing resources, databases, and experimental facilities. One major goal would be the creation of a national knowledge bank. Such a knowledge bank would capture and organize computer-based knowledge in various scientific fields that is currently available only in written/printed form, or in the minds of experts or experienced workers in the field. This knowledge would be stored in knowledge banks that would be accessible over the network to individual researchers and their programs. The result will be a codification of scientific understanding and technical know-how in a series of knowledge-based systems that would become increasingly capable over time.

## **16.6. CONCLUSION**

In this paper, we have tried to describe the functions required of the interconnected national network to support scientific research. These functions range from basic connectivity through to the provision for powerful distributed user services. Members of the Internet Task Force on Scientific Computing are listed in Section 16.7.

Many of the goals described in this paper are achievable with current technology. They require coordination of the various networking activities, agreement to share costs and technologies, and agreement to use common protocols and standards in the provision of those functions. Other goals require further research, where the coordination of the efforts and sharing of results will be key to making those results available to the scientific user.

For these reasons, we welcome the initiative represented by this workshop to have the government agencies join forces in providing the best network facilities possible in support of scientific research.

## 16.7. INTERNET TASK FORCE ON SCIENTIFIC COMPUTING

Rick Adrion  
University of Massachusetts

Ron Bailey  
NASA Ames Research Center

Rick Bogart  
Stanford University

Bob Brown  
RIACS

Dave Farber  
University of Delaware

Alan Katz  
USC Information Science Institute

Jim Leighton  
Lawrence Livermore Laboratories

Keith Lantz  
Stanford University

Barry Leiner (Chairman)  
RIACS

Milo Medin  
NASA Ames Research Center

Mike Muuss  
US Army Ballistics Research Laboratory

Harvey Newman  
California Institute of Technology

David Roode  
Intellicorp

Ari Ollikainen  
General Electric

Peter Shames  
Space Telescope Science Institute

Phil Scherrer  
Stanford University

## 17. SUMMARY OF AGENCY NETWORK PLANS AND USER COMMUNITY

---

**Stan Ruttenberg**  
Computer Science Network

### 17.1. NATIONAL SCIENCE FOUNDATION AFFILIATED NETWORKS

#### 17.1.1. NSFnet Backbone Network and Connections to the Research Community

Currently, the National Science Foundation (NSF) network consists of 56-kbit/s terrestrial lines that interconnect the NSF supercomputer centers. Network management and technical services are provided by University of Illinois and Cornell University. User services (communications) are supervised by the University Corporation for Atmospheric Research (UCAR) and operated by Bolt, Beranek and Newman under subcontract. Advanced software and communications protocol are being developed by the Information Sciences Institute at the University of Southern California.

Access to the NSFnet backbone is available to researchers in a variety of ways, including

- a consortia managed by three of the supercomputer centers (San Diego, Princeton, NCAR),
- ARPANET connections,
- some CSNET and BITNET connections, and
- by any terminal that can access Telnet.

To expand the system nationwide, additional access routes are now being developed through state or regional networks; these networks are being organized by local groups with start-up funding by NSF. At this time, such networks are coming into operation in such areas as New York, Michigan (an existing network that is being modified to provide IP access), the mid-Atlantic states, Texas, the upper Mississippi area of the Big Ten universities, the Pacific Northwest, the Rocky Mountain region, and the San Francisco Bay region. Other networks are in planning stages and it is anticipated that by the end of 1987 most of the nation will be covered by such networks, which will provide access to perhaps 200 university campuses.

#### 17.1.2. NCAR Supercomputer Center

The National Center for Atmospheric Research (NCAR) Supercomputer Center is operated by the University Corporation for Atmospheric Research (UCAR) under contract with NSF. Since their inception, NCAR's computers have been used by the university community as well as by resident NCAR scientists. At present, use of the center is roughly divided between 50% university use and 50% nonuniversity use, although some of the outside use actually involves collaborative efforts with NCAR scientists.

### 17.1.3. Remote Job Entry Capability

#### 17.1.3.1. UNINET

When the NCAR mainframe became large enough to permit remote use without giving up much of the CPU efficiency, batch remote job entry (RJE) capability was initiated around 1972. In 1982, a 1.2-kbit/s packet switched leased-lines system (UNINET) was introduced. Usage of UNINET grew rapidly and it is now a fixed component of the NCAR computer access system. Many small users keep their files (programs and data banks) at NCAR and submit execution instructions over the RJE facility. Limited output, including some graphics for preview use, is available back to users via the leased lines, but large outputs are usually sent back via mail. Approximately 150 institutions are connected, making up a total of perhaps 500-1,000 users of this facility. Although installation of 2.4- or 4.8-kbit/s lines might be adequate for the RJE utility, a few users need more bandwidth for input and/or output, so a few 9.6-kbit/s lines were installed. Later, a satellite-link system was developed and funded by NSF on a pilot basis.

The NCAR computer facility has surveyed its user community several times. Reference 1 shows the trend of requirements of our community as of 1983. The questions asked in this survey could serve as a guideline to an updated national survey, which might be a useful task for this group to undertake in preparation for the long-term projection requested by the Congress.

#### 17.1.3.2. USAN

In 1986, NCAR began implementation of a star system called the University Satellite Network (USAN) using satellite video subcarrier capability. The central hub at NCAR has a 224-kbit/s capability, and five university sites have a receive/transmit capability at 56 kbit/s. Additional sites may be added. The original site selection criteria included large users, geographic distribution (to test reliability at low antenna elevation angles), and varied climatic conditions (dry, rainy, icing, etc.). This network is now operational in a pilot mode. Modelers running large experiments (for example, at the University of Miami) are obtaining satisfactory supercomputer service from NCAR.

A zero-order estimate of possible future uses for the USAN includes

- 15-30 institutions where large-scale atmospheric/oceanographic computing is conducted, and
- perhaps 50-100 modelers running complex experiments.

Input of 56 kbit/s to NCAR may remain satisfactory for a short time, but some users will need 224-kbit/s or larger capability eventually. A return path of 56 kbit/s to the user site will soon become unsatisfactory because output files will be too large to facilitate the final data processing at the home site. Since it is desirable to offload large data-processing jobs from the host supercomputer, high-capability return paths seem mandatory in the near future. Moreover, there is interest in online interactive graphics capability by scientists running large and complex model experiments. This would also require higher bandwidth than available at present for the return path.

#### 17.1.3.3. UNIDATA

The University Data (UNIDATA) project was initiated primarily to fill the data distribution gap when the National Weather Service (NWS) and the National Oceanic and Atmospheric Administration (NOAA) decided to deactivate their teletype data distribution to universities and private sector. In addition, NSF and UCAR decided to try to develop community "universal" local

software/communications systems so that small and large schools could implement departmental data reception, archival (as required) functions, data processing for research and education, and extensive local file-server and advanced workstation capabilities.

NWS data (19.2 kbit/s) are accessed at Washington, DC, uplinked at Chicago by a commercial service, and broadcast to TVRO earth stations at 40-100 universities. An additional 19.2-kbit/s bandwidth is available for uplinking special data sets produced at some advanced university analysis centers. The present available bandwidth is barely adequate for high-resolution, sectorized satellite images, and higher bandwidth will be needed soon if universities are to be able to make use of all the advanced meteorological analysis products now becoming available. NWS (see section below on NOAA) also needs to increase its capability to distribute normal weather data to its field stations as well as to distribute the high-resolution satellite data now available. These needs will grow by orders of magnitude in the next 10 years.

In the future, meteorological departments will want to access each other's specialized data sets, and request and receive data sets from NCAR and the NOAA archives using the UNIDATA system. Geophysical scientists also will want to collaborate using remote logon facilities and will want to exchange software via some national communication system. Electronic mail is now a way of life among geophysical scientists, so an electronic mail service totally integrated with a data network will be desirable. However, because UNIDATA was not created as a mail system, reliance will be on NSFnet. Thus, the meteorological/oceanographic community will become a large user of NSFnet.

Because meteorologists conduct extensive field experiments, a way is needed to communicate data to the field operations center, process them rapidly, and redistribute some selected, high-resolution products back to such field operations as aircraft control or radar centers, so that aircraft and radar operations will be able to vector in on the fast-moving meteorological systems under study. In addition, some near real-time products might have to be sent to local experimental forecast centers, not only as part of the field experiments, but also to serve as training aids for forecasters. It may also be important to send such information to some university departments for additional analysis, for consultations on the progress of the experiments, and for pedagogical purposes. A national research communications network with high-bandwidth capabilities, supplemented by specialized mobile communication systems for the temporary field sites, may supply part of the backbone communications for such uses.

#### 17.1.3.4. CSNET

The Computer Science Network (CSNET) was founded by an NSF in 1981. Some 200 sites are connected to CSNET in a variety of services. As of November 1986, the activities conducted on CSNET included the following.

- There are 141 PhoneNet sites connected via ASCII electronic mail. The sites are served primarily by 1.2-kbit/s packet switched lines and some 2.4-kbit/s lines. Local hosts are interrogated each night, resulting generally in 1- to 2-day mail service. Sites that desire faster service call in directly to the CSNET relay in Boston. A major goal of CSNET is to upgrade this service to Internet Protocol (IP) capability for full file-transfer and near real-time mail service.
- The 16 X.25 sites have full IP connectivity to the internet.
- The 33 Advanced Research Projects Agency (ARPA) sites have full connectivity to the internet.
- There are four CYPRESS sites. CYPRESS is an experimental network with IP connectivity through 9.6-kbit/s leased lines, using specially programmed minicomputers (for example, MicroVax II and Sun).

CSNET hosts at university sites are generally housed in the computer science departments and not available campus-wide, though this situation is now gradually changing as a result of NSFnet policies. CSNET is moving towards an extension of computer science hosts to campus-wide hosts.

CSNET is undertaking a development program, exemplified by CYPRESS, to extend IP connectivity to small schools and to clusters of industrial sites. Also, CSNET operates Nameserver and Info Server facilities that may be developed further for wider use in NSFnet.

#### 17.1.4. International Gateway Usage

International gateways are located in Australia, Canada, Finland, France, FRG, Israel, Japan, Korea, New Zealand, Sweden, Switzerland, and the United Kingdom.

The following chart shows international gateway usage during the period January 11 (6 AM) to January 18 (6 AM). Information is presented as follows:

Numbers of messages (Count)

Number of addressees (Addrs)

Numbers of characters (Chars)

Mail sent to the site (OUT)

Mail received from the site (IN)

Number of completed PhoneNet calls (OKs)

Number of incomplete calls noticed (NOs)

Total connect time (Seconds)

Some recent traffic statistics are shown in this table:

Name	OUT			IN			CALLS		Seconds
	Count	Addr	Chars	Count	Addr	Chars	OKs	NOs	
Australia	86	88	159990	40	40	57686	30	4	6980
Canada	1217	1288	5651539	338	343	622540	174	1	87700
France	3	5	3019	0	0	0	1	0	211
Germany	486	490	2306554	54	54	59376	7	7	29315
Israel	100	101	130530	68	68	60210	43	2	8476
Japan	257	261	827411	208	202	433474	73	2	27829
Korea	82	84	255668	63	67	92027	8	2	5779
Sweden	116	125	678602	19	18	18752	46	2	13925
Switzerland	69	70	398214	58	57	51018	142	10	12024
TOTAL	2416	2512	10411527	848	849	1395083	524	30	192239
TOTAL	18572	21223	82573560	7014	7456	18390828	2699	355	1810452

(All PhoneNet)

Note: Traffic for French gateway may be understated because a new gateway operated by INRIA was being established during this period.



## **17.2. DEPARTMENT OF ENERGY (DOE)**

### **17.2.1. Energy Sciences Network**

The DOE's Energy Sciences Network (ESnet) has one CRAY-2, one CRAY X-MP/22, and two CRAY-1 supercomputers. Access is via the Magnetic Fusion Network (MFEnet).

The MFEnet is divided into the MFE I and MFE II. MFE I currently oversees MFE protocols and closed access. MFE II oversees moves from closed to open access, migration to IP and X.25 access and use of multiple protocol suites, additions of new communities, establishment of gateways to Japan, FRG, Switzerland (CERN), additional selected performance enhancements, and electronic mail gateways. Functions include remote terminal access file transfer, electronic mail, various servers and naming conventions, gathering of statistics, and access to supercomputers.

The near-term goals of MFE II include extending bandwidth from 56 kbit/s to 1 Mbit/s, completing terrestrial interconnectivity, establishing International Standards Organization (ISO) IP routing, extending capabilities to the home environment (e.g., PCs), using TCP/NSP and TCP/IP on Cray supercomputers, and using ISO protocols and further extension of multiple protocol suites. New functions envisioned will include specialized distributed processing, remote procedure calls, task-to-task communications, type-of-service routing, PC hosts, and workstation support.

### **17.2.2. High-Energy Physics Net (HEPnet)**

In the high-energy physics world, connectivity is contemplated with about 10 major laboratories around the world, and 100-400 scientists in 10-40 institutions in 12 countries through X.25 services.

HEPnet has a backbone 556-kbit/s trunk that services 9.6-kbit/s feeder lines, with one satellite link to CERN and one to Japan. By FY 1989, the trunk will be doubled, a 56-kbit/s line over transatlantic fiber will be established to CERN, and migration will begin to ISO standards.

## **17.3. NATIONAL AERONAUTICS AND SPACE ADMINISTRATION (NASA)**

### **17.3.1. NASA Aerodynamics Simulation (NAS)**

The initial configuration of the NAS network system has been composed of a HSP-1 computer that has 250 MFLOPS and 256M 64-bit words in its central memory. Lines to remote sites transmit at 56 kbit/s, and lines between relevant NASA centers transmit at 1.544 Mbit/s. Plans are under way to upgrade the computer to 1 GFLOPS with 6.2-Mbit/s lines to centers. Future options include upgrading the computer by a factor of 4 to 5.

### **17.3.2. Program Support Communication Network (PSCN)**

The network backbone for the PSCN has been contracted with Boeing and RCA. Present capabilities include T1 terrestrial and satellite links, data supply, facsimile, teletype and voice circuits, video teleconferencing, and batch file transfers.

PSCN associated facilities include a NASA packet switched subsystem with 9.6-kbit/s synchronous and asynchronous lines to NASA researchers; a computer network subsystem that produces high-speed file transfers between NASA supercomputers, as well as having store and forward system capabilities;

and a space physics analysis network with a 56-kbit/s backbone with TCP/IP protocols, and links to the DDN and NSFnet, to NASA PIs, and to the European Space Agency, BITNET, TELENET, and ARPANET.

#### **17.4. NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION (NOAA)**

The major communication needs in relation to NOAA are in the NWS and the National Environmental Satellite and Data Information Services (NESDIS). Most of this need is operational. Routine operational data are distributed to the NWS forecasting centers (about 50) by a closed-circuit interactive system, but additional capability is needed to transmit high-resolution satellite images to these centers and to distribute selected analysis products to the 200 other local forecasting stations.

NOAA networking also encompasses major research laboratories, which need online data and access to data archives. In addition, routine and specialized observations from NWS and NESDIS services are widely used for value-added analyses by the commercial sector and by universities for research and teaching.

NOAA has plans for two systems to be implemented in the near future. The first, NOAAPORT, will be a near real-time distribution of weather and environmental data to users. In 1987, this system will use the DOMSAT capability to distribute sectorized images from the geostationary satellites and selected analysis products to a limited number of forecast sites. It is expected that NOAAPORT will be fully implemented by 1991.

The second system, NOAAANET, will support retrospective data management and user services, including research. Pilot operation is planned for 1987 and will encompass some data management and demonstration projects of user services. Depending on the experience gained from the pilot operation, plans will be formulated for an operational system to be implemented in FY 1991.

#### **17.5. U.S. NAVY**

Major Agency network plans for the US Navy are currently under way at the Institute of Naval Oceanography (INO). A supercomputer at INO (Bay St. Louis, Mississippi) may be connected to NCAR via a USAN link, which would, in effect, connect INO to the NSFnet backbone.

At the Navy Research Laboratory, a CRAY X-MP/12 is now in operation with 36 1.2-kbit/s and 9 4.8-kbit/s dialin lines. Three 56-kbit/s lines are available to the DDN/MILNET. Future plans may include a link to NSFnet, perhaps via a connection to the INO-NCAR USAN link.

#### **17.6. NATIONAL INSTITUTES OF HEALTH (NIH)**

Currently at the NIH, networking capability consists of a CRAY X-MP/22 linked to some 4,000 users working in molecular modeling.

## 18. DOE NETWORKING REQUIREMENTS

---

**D. F. Stevens**  
Lawrence Berkeley Laboratory

### 18.1. THE DOE RESEARCH ENVIRONMENT

The Department of Energy (DOE) consists of more than 50 government and contractor installations, plus a considerable number of ORUs (Organized Research Units), and individual researchers. (A list of the principal DOE research sites is provided in Section 18.6) The DOE scientific research community is distributed across the entire establishment and comprises a total of about 30,000 research scientists. The geographic distribution includes all 50 states, plus collaborators in Canada, Europe, and Japan. These scientists all require access to computers, both locally and at centrally operated national or international facilities. Approximately 3,000 computers,\* ranging in size from minis to the largest available supercomputers, are operated within the 50 DOE sites. Most sites have installed, or are developing, extensive local area networks for internal communications.

One characteristic of DOE research has been the development of *national* facilities for use by the entire scientific community. All of DOE's major accelerators have been operated in this manner, as well as such diverse facilities as the Atomic Resolution Microscope at Lawrence Berkeley Laboratory (LBL) and the supercomputers at the National Magnetic Fusion Energy Computing Center (NMFEC) and Florida State University (FSU). As a result of this tradition, DOE research has come to include very large national (and *international*) collaborations, involving hundreds of individuals from many institutions.\*\* In such an environment, rapid and effective communication (of messages, extended text, data, and graphics) is not a luxury but a necessity. A recent survey conducted by DOE's Office of Energy Research indicates that, within each major DOE site, an aggregate of thousands of messages and files are transmitted and hundreds of remote terminal hours are spent daily.

### 18.2. EXISTING DOE NETWORK USAGE

In order to accomplish their work, research scientists within DOE are today required to use a large number of different networks and ad hoc communication links. The principal channels include MFEnet (to support use of the supercomputers at NMFEC and FSU by the Magnetic Fusion Energy Program and the Office of Energy Research Supercomputing Program), HEPnet (for the High Energy Physics community), BITNET, ARPANET/MILNET, TYMNET, and a large number of point-to-point leased lines (more than 100 in the High Energy Physics community alone). MFEnet and HEPnet are DOE-developed networks; BITNET is sponsored by essentially the whole of the US and European academic communities; ARPANET and TYMNET are commercial (quasi-commercial, in the case of ARPANET) services.

External networking is used extensively by DOE scientists to support three types of access to distant resources: remote interactive access to computers or information services at other sites, the transfer of large blocks of data from one site to another, and message-oriented services.

---

\*This number does not include PCs, desktop publishing systems, CAD/CAE/CIM stations, or other intelligent workstations. It also does not include computers in universities to which DOE-sponsored scientists may have access.

\*\*Current collaborations in high energy physics experiments typically consist of 100-400 scientists representing 10-40 institutions from a dozen countries.

1. Remote interactive access: There are two general situations in which a scientist seeks remote computing access to another site:
  - to achieve access to a computer with capability beyond that which is available locally, or
  - to achieve access to data that, because of currency or quantity, cannot be transmitted from their place of residence to that of the user. In the former case, remote access replaces acquisition; in the latter, it replaces travel.
2. File transfer: This is in some sense a generalization of a message service, where the "message" is a program or data file. As a result, the volume of data is two or three orders of magnitude greater than for a message-oriented service, and the destination is often a program or system utility (PRINT, for example) rather than a person.
3. Message-oriented services:
  - a) Person-to-person messages (electronic mail): This type of service sees heavy use on all networks. It is the principal reason for the widespread implementation of BITNET within DOE. Many members of the extended energy research community are accessible (for electronic communication) through no other means.
  - b) Bulletin board services: The amount and effect of bulletin board traffic on networks is easy to underestimate. Users tend to see bulletin boards as a *local* service and fail to realize that essentially all bulletin board messages either originate off-site or have at least one off-site addressee. Bulletin boards exist at all levels of specificity and practicality. The most common usage is to transmit system news about the host site, but there are also discipline-specific bulletin boards, and an increasing number of useful commercial bulletin boards (e.g., Autocad, Byte Magazine bulletin board).
  - c) Electronic conferences: These are generalizations of both bulletin board services and electronic mail. They are more interactive than most bulletin boards and have longer memories than either bulletin boards or electronic mail services. They contain utilities to assist users to look up old submissions, to vote, to contribute, to engage in private dialogue with other conferees, etc.

The estimates given below in Section 18.4 are limited to remote access and file and message traffic. No attempt has been made to address the traffic that would be generated by newer developments such as real-time video conferencing.

### 18.3. FUNDAMENTAL REQUIREMENTS

There are three fundamental requirements for a network:

1. Connectivity. Every site needs access to every other site. (We include adequate availability and reliability in this requirement.)
2. Bandwidth. Will the connection support the traffic I generate?
3. Speed. For some applications, particularly those involving direct terminal access to a remote site, raw bandwidth is not enough; the connection must also provide a response time that approximates that available to a local connection.

Not every connection needs to support the maximum bandwidth with the minimum response time, of course. It is probable that most of them need to support only message traffic and occasional file

transfers of moderate size (up to the order of 100 kbyte), but there are some that need the capability to pass massive amounts of data hourly (of the order of tens of gigabytes), some that need to provide sub-second response time for interactive access, and a few that need to do both.

## 18.4. DATA BANDWIDTH REQUIRED

The required\* bandwidth\*\* estimates below address only traditional types of network traffic; there has been no attempt to allow for such uses as remote experimental control (such as might be necessary to manage and experiment in a satellite or space station) or extensive real-time video conferencing. The estimates are based upon general experience rather than extensive surveys. They are thought to be moderately conservative in the sense that they are adequate to meet the needs of the DOE research community for the next 2 to 3 years.

### 1. Human-generated material:

- a) Message: 10 messages x 10 kbyte/message x 10 (average distribution list)  
= 1 Mbyte/person/day
- b) File: 10 text pages/day + 1 graph + 10 remote program submissions @ 5 kbyte  
= 30-kbyte text + 1-Mbyte graphics + 50-kbyte program  
= 1.1 Mbyte/person/day
- c) Terminal Access (text): 1 hour/person/day @ 1 keystroke/second and 45 byte/keystroke  
= 3600 x 1 x 45 = 160 kbyte/person/day
- d) Terminal Access (graphic): 10 full screens/person/day @ 1 Mbyte/screen  
= 10 Mbyte/person/day

Total human-generated material is thus 12.26 Mbyte/person/day.

### 2. Machine-generated material:

- a) Data: 1 x 6250 bpi tape/person/day = 200 Mbyte/person/day
- b) Output: 250 pages/job = 250 x 10 x 10K/page = 25 Mbyte/person/day

Total machine-generated material is 225 Mbyte/person/day.

- 3. The total *average* data requirement is thus 237 Mbyte/person/day, of which perhaps 90% would be transmitted during a 10-hour prime-time period, for a prime-time average of 5.9 kbyte/second/person. It must be emphasized that this number reflects data traffic only, and is an average over the whole of prime-time. Peak traffic (averaged over the busy minute) characteristically runs a factor of 10 higher than average traffic (if the network has sufficient capacity to sustain such a load).

---

\*The potential "required" bandwidth is infinite unless we learn how to control the expansion and proliferation of distribution lists. Electronic systems do not just replace paper systems on a one-for-one basis. Because it is far easier to make and send (and forward) copies in an electronic system, it is done; and because its distribution lists are easy to extend (by appending other lists) and difficult to prune, they continue to grow without bound.

\*\*"Data bandwidth" means user data plus envelope (addressing information); it does not include any protocol bits, necessary retransmissions, or any other system-induced traffic. Such considerations could increase the requirements by anywhere from 10% to 100%.

4. **Distribution:** The traffic is not evenly distributed. The bulk (perhaps 50-70%) is concentrated along a corridor linking Lawrence Berkeley Laboratory/Stanford Linear Accelerator Center/Lawrence Livermore National Laboratory on the West Coast with the Argonne and Fermi National Laboratories in the Chicago area, and continuing on to Brookhaven, Princeton, and Massachusetts Institute of Technology on the East Coast. There is also a major southern arc including Los Angeles, Boulder, Los Alamos, Austin, Tallahassee, and Oak Ridge, and extensions into all parts of the US, plus Japan, Canada, and Europe. (See the network maps, Figures 1-6, in Section 18.7).

## 18.5. CURRENT PLANS AND PROJECTIONS

A glance at the maps in Section 18.7 will show the existence of a great deal of redundancy of routing among the networks used by DOE. To some extent, this is unavoidable, because the networks are provided by independent suppliers (ARPA: DCA; BITNET; the universities; TYMNET: TYMNET Inc.; etc.). It can also be seen that MFEnet and HEPnet currently provide duplicate paths in some places so that the necessary bandwidth can be achieved. DOE has begun to merge the two networks, under the name of ESnet (Energy Science Network), and it is expected they will have become a single network by the end of 1990. (A map of the first stage of ESnet, to be completed in 1987, is included in Section 18.7).

The network traffic of the future can be split into two parts: a portion that can sustain the delays inherent in satellite transmission and a portion that cannot. For most purposes, we can consider the second class as consisting of terminal access: There are very few terminal applications that can sustain the 0.4-s additional response-time delay resulting from satellite use (and few other applications that cannot). The bulk of the traffic, therefore, is, and will remain, satellite-amenable. It will also be the slowly growing component of future network usage. Its growth will be determined by growth in the scientific research establishment, compounded by growth in the complexity of the research undertaken, but is unlikely to exceed 10% per year.

The satellite-unsuitable portion of the traffic will grow explosively, however, as scientific computing makes more use of the possibility of remotely driven, supercomputer-generated, interactive graphics. We are here certainly below the rising knee of a classic learning curve, and can expect growth to accelerate from perhaps 10% per year at present to 50% per year for several years somewhere around 1993. This could result in traffic demands of 20 Mbyte/person/day by 1990 and 120 Mbyte/person/day by 1995 (Figures 7-10). That appears to be a rather modest requirement, in view of today's total requirement of 237 Mbyte/person/day, *but it is an interactive requirement*. A single screen, including color and/or gray-scale, will require 10 Mbyte instead of only 1 Mbyte, but with a response time no worse than half a second. In other words, each online scientist will need instantaneous access of 20 Mbyte/s to the computer of his choice.

## 18.6. PARTIAL LIST OF DOE RESEARCH SITES

(Includes foreign collaborations)

Ames Laboratory	Cornell University
Argonne National Laboratory	University of Florida
Bettis Atomic Power Laboratory	Florida State University
Brookhaven National Laboratory	Harvard University
Battelle Project Management Division	University of Houston
Fermi National Accelerator Laboratory	University of Illinois
GA Technologies, Inc.	University of Indiana
Grand Junction Project Office	Johns Hopkins University
Idaho National Engineering Laboratory	Michigan State University
Inhalation Toxicology Research Institute	Purdue University
Knolls Atomic Power Laboratory	Stanford University
Lawrence Berkeley Laboratory	SUNY, Stony Brook
Lawrence Livermore National Laboratory	University of Texas
Los Alamos National Laboratory	University of Washington
MIT Laboratory for Nuclear Science	University of Wisconsin
National Magnetic Fusion Energy Computing Center	Yale University
NYU Courant Mathematics and Computing Laboratory	University of Bonn (FRG)
Oak Ridge Associated Universities	DESY (FRG)
Princeton Plasma Physics Laboratory	Heidelberg University (FRG)
Sandia National Laboratories, Albuquerque	MPI-Munich (FRG)
Sandia National Laboratories, Livermore	Weizmann Institute (Israel)
Solar Energy Research Institute	University of Tel Aviv (Israel)
Stanford Linear Accelerator Center	Technion University (Israel)
Waste Isolation Pilot Plant	University of Bari (Italy)
Boston University	University of Bologna (Italy)
Brandeis University	University of Frascati (Italy)
University of British Columbia	University of Genova (Italy)
Brown University	University of Milano (Italy)
California Institute of Technology	University of Padova (Italy)
Carnegie-Mellon University	University of Roma (Italy)
UC Berkeley	University of Trieste (Italy)
UC Davis	KEK, University of Tokyo (Japan)
UCLA	University of Tsukuba (Japan)
UC San Diego	Catholic University of Nijmegen (NL)
UC Santa Barbara	CERN (Switzerland)
University of Chicago	University of Geneva (Switzerland)
University of Cincinnati	Birmingham University (UK)
University of Colorado	Bristol University (UK)
Colorado State University	Cambridge University (UK)
Columbia University	

### 18.7. NETWORK MAPS AND PROJECTED GROWTH DATA

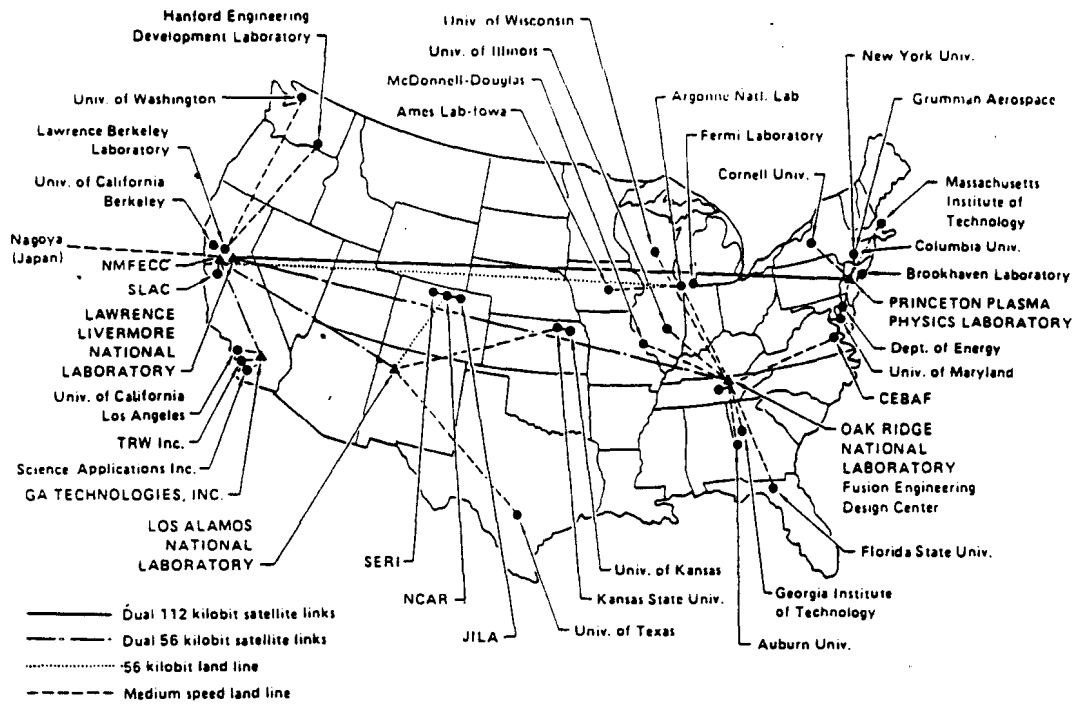


Figure 1. MFEnet.

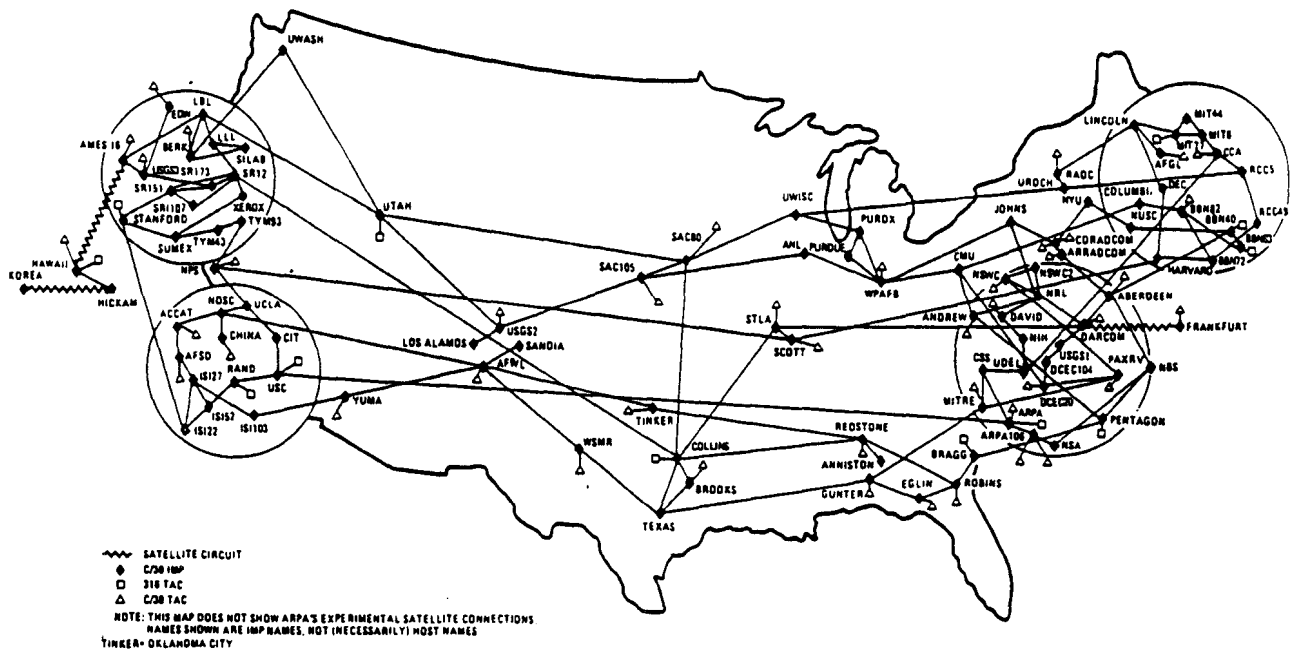


Figure 2. ARPANET.



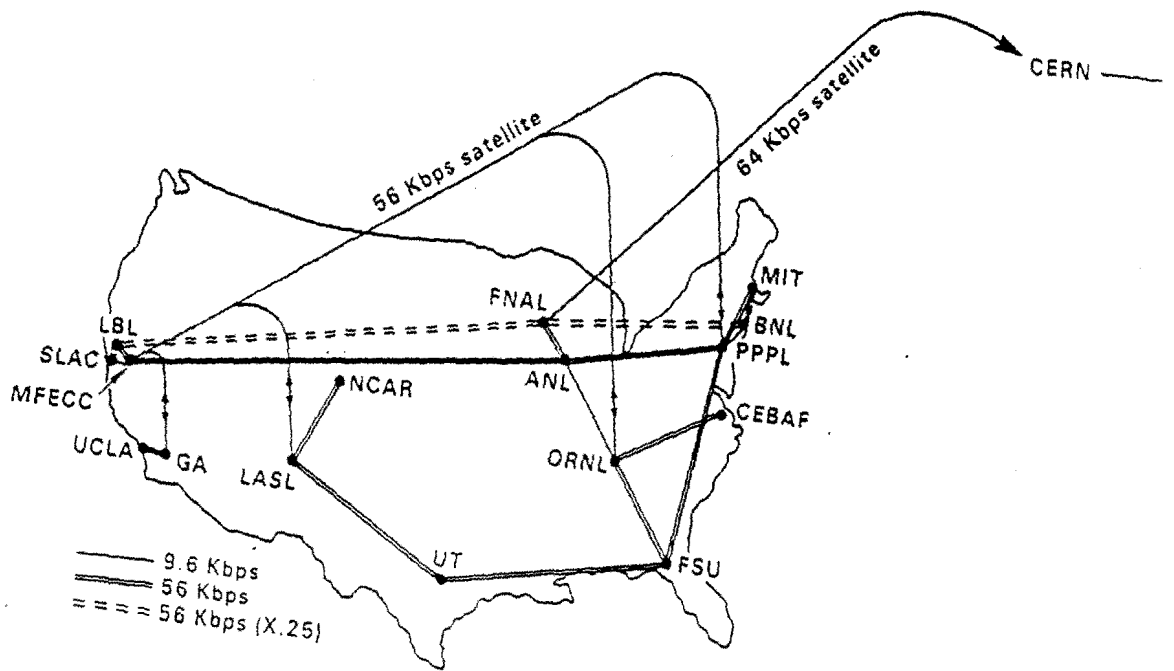


Figure 3. ESnet, Stage 1.

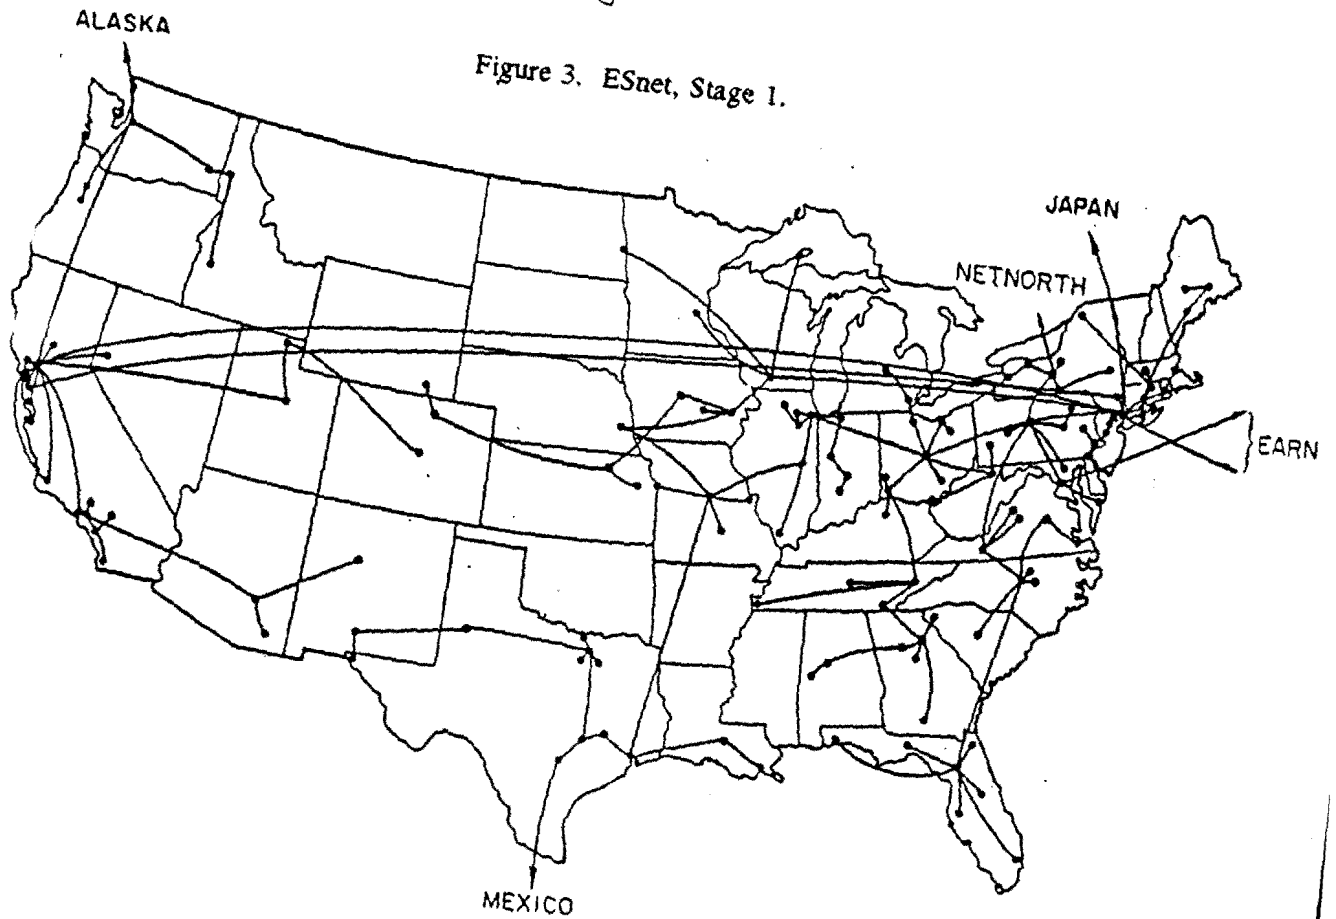


Figure 4. BITNET.




Figure 5. DECnet component of HEPnet (PHSYNET).




Figure 6. HEP leased lines for terminals.





Figure 7. Estimated DOE quick-response network bandwidth required as a percentage of total bandwidth.

Figure 8. Estimated DOE quick-response network bandwidth required in Mbytes/person/day.







Figure 9. Estimated DOE total network bandwidth required in Mbytes/person/day.

Figure 10. Estimated DOE quick-response network bandwidth required in Mbytes/person/day.



## 19. RESEARCH UNIVERSITY BANDWIDTH ESTIMATES

---

**Glen Ricart**  
University of Maryland

To estimate the bandwidth needed to support research at a major research university, the kinds and volume of traffic that would be used were evaluated assuming the networking resources were available in a 10-year time frame.

There are three somewhat overlapping views of the amount of bandwidth needed for the network:

- Bandwidth needed by type of task
- Information needs by type of user
- Information flowing at the university boundary

Numbers were developed for each model so that intelligent guesses could be made about the University of Maryland traffic needs in 10 years when not constrained by current technology or cost.

In each model, we estimate the peak bandwidth needed for each type of service. For example, in the Task Model, the campus bandwidth needed is dominated by the need for at least one campus researcher to receive full color and full motion high-resolution images. These images may be real-time three-dimensional visualizations of real objects such as the human body under nuclear magnetic resonance (NMR) scans or mechanical parts undergoing stresses, or they may be simulations of the dynamic behavior of these objects. More than one researcher on campus may need access to this kind of high-quality image, but the bandwidth required is developed on the assumption that such a high-quality link can be serially reused. Similar assumptions underlie the other estimations.

In the following tables we have made some assumptions:

- In the User Type Model, we estimate typical bandwidth needed based on user type.
- In the Edge of Campus Model, we examine data flow across the border of the campus by source and destination.

By any of these estimation techniques, the University of Maryland needs about a gigabit per second data pipe to realize the research gains possible with a national academic network. While these are extremely rough estimations, it is likely that other major research universities will have similar requirements.

TASK MODEL						
Task Description	Ratio of Burst Size to Switch	Bandwidth Needed for One User	Number of Users	% Time on Task	Peak User-Count	Bandwidth Needed Overall
Full motion <sup>1</sup>	very high	2 Gbit/s	4	10%	1	2 Gbit/s
Color images	high	100 Mbit/s	40	10%	4	0.4 Gbit/s
B&W images	moderate	4 Mbit/s	100	15%	15	0.06 Gbit/s
RPC	low	1 Mbit/s	1000	10%	100	0.1 Gbit/s
Data retrieve	moderate	1 Mbit/s	1500	5%	75	0.07 Gbit/s
Graphics	low	100 kbit/s	2000	5%	100	0.01 Gbit/s
E-mail	low	1 kbit/s	2500	15%	375	0.00 Gbit/s
Virtual Term	very low	10 kbit/s	2000	5%	100	0.00 Gbit/s
Total						2.64 Gbit/s <sup>2</sup>

<sup>1</sup> 2K by 2K pixel resolution at 24 frames per second with 8 bits for each of three colors.

<sup>2</sup> The 2.64 Gbit/s estimate does not include protocol overhead. Current protocols need a line with 100% greater capacity to include protocol overhead and unevenness of data flow.

USER TYPE MODEL					
Type of User	Number on Campus	Typical bandwidth Need	% Time on Task	Peak Count	Bandwidth Needed
Undergrads	30,000	56 kbit/s	20%	6,000	336 Mbit/s
Graduate	5,000	112 kbit/s	40%	2,000	224 Mbit/s
Faculty	1,500	1 Mbit/s	30%	450	450 Mbit/s
Research staff	1,500	112 kbit/s	50%	750	84 Mbit/s
Total					1.09 Gbit/s

Note: The total is in the gigabyte per second range, but is smaller than the previous estimate because the "typical" user model employed here does not take account of the usual high-resolution full motion user of the task model.

EDGE OF CAMPUS MODEL			
Source/Destination	Users	Number of Rate	Data Needed
Remote real-time experiments (for example, CEBAF)	2	200 Mbit/s	400 Mbit/s
Supercomputer centers	60	1 Mbit/s	60 Mbit/s
Imaging/animation centers	50	10 Mbit/s	500 Mbit/s
Other universities	500	100 kbit/s	50 Mbit/s
Total			1.01 Gbit/s

## GOVERNMENT ROLE





## 20. THE FEDERAL GOVERNMENT'S ROLE IN NATIONAL RESEARCH NETWORKING

---

Sidney Karin, Susan Estrada, and Daniel Drobnis  
San Diego Supercomputer Center

### Abstract

*The federal government's role in networking should be oriented toward setting goals for and establishing (and in many cases furnishing) stable funding appropriate to network and project types. Management and funding should be kept separate. Networking projects, both new and existing, should be separated into production networks, oriented toward users; experimental networks, oriented toward new network development; and network research projects, oriented toward research on specific networking and communications technologies. A permanent advisory group should be established, through the Federal Coordinating Council on Science, Engineering and Technology committee, with broad representation from network researchers, communications technologists, and most particularly, scientific- and engineering-oriented network users. This group should be primarily concerned with seeing that projects are properly categorized, definite goals are established, and progress toward those goals is monitored and reported.*

### 20.1. DEFINITIONS, FUNCTIONS, AND GOALS

Several types of research networks and projects are discussed in this paper, and it is important to understand their unique characteristics. So a few words of definition are in order.

The term production network is used to refer to a network, the majority of whose users are doing research on topics not related to communications or computing. It needs to be easily accessible, designed for high availability and reliability, and optimized for the functions such as file transfer or interactive computer access, which most of its users require.

An experimental network might also be used by noncommunications researchers, but would be primarily a testbed for new ideas and technologies, or optimized for special functions. Because of this, availability and reliability of an experimental network would not be expected to be as high as for a production network.

Network research projects would, as the name implies, include investigations aimed at developing new technologies in areas such as communications or distributed processing, and would not be expected to function in a production environment until their ideas have been proved and refined.

The success of production networks hinges upon several factors. First, the services provided must be those perceived by the user community in question to be the ones required. Network providers must recognize the importance of these perceptions. Provision only of services deemed to be necessary by network designers may preclude success if user perceptions are not understood and taken into account. Conversely, the provision of services for which the user community has no need expends scarce resources without purpose.

The success of experimental networks hinges upon different factors; the very definition of success is different. Whereas success for a production network means productive use by the intended community for the intended purpose, success for an experimental network may mean a transition to a production network status or it may mean a simple proof or demonstration of principle. An experimental network may be judged successful if it serves as a testbed for network researchers to try new ideas and techniques.

Performance and reliability criteria are often less important for experimental networks than for production networks. Funding requirements for experimental networks may be lower simply because they may be implemented on a smaller scale.

In this regard, it is important to note the competing demands of connectivity and performance in computer networks. Network designers have in recent years tended to emphasize connectivity, while network users, once connected, prize performance and reliability.

Performance and reliability are such important criteria to network users that they will permanently abandon use of a facility (networks are only one example) after just one or two experiences of inadequate or unreliable performance. It should be noted that network performance may or may not be improved (at least not cost effectively) by simply increasing nominal bandwidth. Examples of networks whose throughput is one or even two orders of magnitude below nominal hardware bandwidth unfortunately abound. And even where packet throughput for file traffic is adequate, failure to provide priority for terminal traffic can render interactive use impractical. In the same vein, reliability may or may not be improved by increasing connectivity.

## 20.2. ASSURING SUCCESS

Happily, a history of successful development of networking techniques has produced a mature technology that is ready to be exploited. This should provide for connection of researchers to a wide variety of target computers, and for such research-oriented computer centers to be connected to each other. The task remains of making effective plans to do this.

While people doing research in the area of network and communications technology have been well represented in planning efforts up to now, people using networks as a tool to support their research in all of the other basic areas of science and engineering have not been. If an example of the under-representation of working noncomputer researchers is necessary, the composition of the present Computer Network Study group can serve as one such example.

It is now time to assure that decisions regarding present and future networks have strong input from the people for whom the work is being done--people in disciplines unrelated to computing and networking who need computer access to accomplish their work. Therefore, future planning for networks to support general scientific and engineering research must be done with effective input from people doing such research. Certainly, experts on networks and communications techniques must be represented as well, but not to the exclusion of real users.

## 20.3. THE FEDERAL ROLE IN NETWORK RESEARCH

The federal government has an important role to perform in the planning and funding of national computer communications networks to support basic research. In addressing this role, however, it is important to stress that we in the United States should use our society's strengths to enhance our lead in technology and the sciences. One of the strongest contributors to this lead in the past has been the

free exchange of ideas and information. Computer communication networks are now emerging as a preeminent tool for doing this.

So it is only common sense that the federal government take steps to promote the productivity of the scientists and engineers who perform the basic research needed to enhance the competitive position of the U.S. in the global economy. And while networks in and of themselves are rightfully of interest to researchers in computer science and communications, it is only through their use by, and impact upon, the much larger community of scientists and engineers in the U.S. that significant federal funding can be justified.

Funding of national research networking should be driven by the requirements of the network users. The ultimate goal is the development of useful tools for noncomputer people. This must be kept firmly in mind, so that the danger of such efforts being regarded as the intellectual property of the developers is avoided. Three types of funding should be considered: total funding, partial subsidies, and nonfederal subsidies.

Total funding for production networks like NSFnet should be identified and made a long-term priority, so that basic researchers can incorporate use of these networks into their own long-term plans. To this end, any uncertainty in such network funding or goals will lead researchers to make independent plans that will result in wasteful duplication and fragmentation of the desired broad-based communications capability.

A useful model in this regard is the effect that a dependable national and local telephone network has in a deregulated communications environment. Individual businesses and building complexes may make their own decisions about internal wiring and equipment, but all make connection to a single national voice communications backbone.

Funding for a national data communications backbone must also be reliable and free to individual researchers, at least, in order to preclude a powerful disincentive to use. In most instances, particularly in the academic research community, incremental charges for network use, apportioned directly to users, will deter use and may again preclude success.

Partial funding should be considered for experimental networks and network research projects. In such cases, the federal government would provide full funding for an appropriate period, perhaps 2 to 5 years, and then reduce funding as a project either becomes self-supporting in a production mode or serves its purpose. This is a good strategy for agencies (e.g., universities, companies, etc.), and has demonstrated applicability to state-wide or even multi-state areas--once a network has been established and the benefits demonstrated. However, it is not an attractive model for individual researchers, or for unproven technologies. The federal government should not expect a full cost recovery from each project.

Combining partial funding with nonfederal subsidies may be an attractive method for leveraging federal funds for network research projects, but this tact lacks stability for experimental network development, and must be carefully planned and use the most proven technology if considered for funding production network operation. When any technical project needs to obtain part of the necessary funding from a source other than the federal government, inevitable conflicts in objectives arise that must be taken into careful account even for limited network research projects, but that can be detrimental or fatal to longer range and broader based networks. In particular, the uncertainty created by only partial federal funding for production networks carries grave risks that users would not base serious research careers upon them, so the benefit of the investment in such networks would be largely lost.

A better form of funding leverage lies in building upon development that already exists. The San Diego Supercomputer Center (SDSC) has built, for example, on behalf of the National Science Foundation (NSF), upon successful investments in supercomputer and communications technology already made by the Departments of Energy and Defense. By doing this, SDSC has become a viable resource for researchers needing remote access to supercomputer power. Just as importantly, SDSC has been at pains to make the results of its own development and integration efforts available to NSF and other computer centers and networks. Funding emphasis and encouragement should be given to projects that use this type of leverage.

If user requirements are such that readily available commercial networking systems could be used to satisfy them, the federal government should encourage that use. However, since commercial networking systems are based on standards, which of necessity lag behind the state of the art by some years, the federal government should expect, demand, and fund state-of-the-art systems for those networks that have requirements that cannot be met by commercial systems. Industry could be helped greatly through the federal government's sponsorship of early users of products. Examples of technologies that have been helped in this way are fiber, data over satellite networks (NMFEC), and the ETA-10.

#### **20.4. MANAGEMENT AND OPERATION OF NATIONAL RESEARCH NETWORKS**

To differentiate clearly between management and operation, it is important to define these terms. Management refers to the directing or supervising of national research networking as a business. Operation refers to the practical application of principles or processes with regard to any national computer communications networks.

The management of national research network development and backbone service should not be centralized in any existing agency, since competing interagency goals exist. Networking management should be coordinated through the Federal Coordinating Council on Science, Engineering and Technology (FCCSET) committee. The FCCSET committee should be used as a funnel for user needs. Once the FCCSET committee has set some management goals through the use of advisory groups and user surveys, the existing agencies would be made responsible for seeing that the strategies are implemented. The existing agencies should not develop these strategies.

Networking should be driven by the intended users. Indeed, networking should not be networking-people driven. Researchers and agencies should dictate network standards based on in-house requirements. Agencies and researchers that don't have an active requirement should not set performance goals. Through the use of steering committees sponsored by FCCSET, users should have a role in defining the planning and priority stages of networking projects. In existing and future networks, users should have a vote regarding the usefulness of the networking service on a regular basis.

Operation of any federal government-sponsored project is complex. Because the same solution or set of rules must apply to everyone, necessary flexibility and practicality can suffer. To minimize these effects on national networking, the operation of production networks is best left to private enterprise. Because the bottom line of a business is profit, private enterprise tends to have a much stronger goal orientation than the government and less abuse is likely to result.

Scholarly research networks, meant for accessing supercomputers and exchanging of information between researchers, should be controlled by a single agency. The functionality of the network should never be jeopardized as a result of interagency conflicts. Interagency cooperation regarding networking must be stressed so that duplication of effort ceases and maximization of resources begins.

The SDSC's communications network is a useful example of a project sponsored by the federal government through the NSF. Although the funding aspect of the project is not optimal, the SDSC has been able to rapidly implement a supercomputer center as well as a substantial computer communications network that extends from the University of Maryland to the University of Hawaii. In addition, the SDSC was among the first of the NSF-sponsored supercomputer centers to implement the standardized high-level protocols of the TCP/IP suite by providing Telnet service to its supercomputer.

Most varieties of network access are supported in a fashion that does not impose communication costs on most individual users. Significantly, the Steering Committee, which provides policy guidance for the center, is composed largely of people with a strong user orientation.

## 20.5. CONCLUSIONS

The FCCSET committee should be active in the formulation of national goals and funding priorities for networking. This national strategy is needed so that all involved agencies and researchers have a clear understanding of what types of networks are and will be available (production, research, and experimental) and what funding is available for each type. Cost sharing should be undertaken principally for local connections; a national user-oriented backbone network to support researchers that is both financially and technically reliable should be a principal objective. Rules by which the players are expected to abide should be clearly stated, and researchers expected to use the resulting networks should be strongly represented in the formulation of these networks.



## 21. THE ROLE OF THE FEDERAL GOVERNMENT IN NATIONAL RESEARCH NETWORK REQUIREMENTS

---

Sidney Karin  
San Diego Supercomputer Center

### Abstract

*The role of the federal government in the management, operation, and funding of national research networking requirements should be as follows:*

- *The federal government should fund and install a high-bandwidth, redundant, backbone network capability for research use.*
- *The federal government should, through established contract and grant mechanisms, make it attractive for research-oriented universities to interface to this backbone network.*
- *The federal government should move quickly to establish an accepted protocol for networks, and make it attractive for regions and states to utilize such protocol.*

### 21.1. BACKGROUND

On June 10, 1985, a Hearing before the Subcommittee on Energy Development and Applications, and the Subcommittee on Science, Research and Technology, of the Committee on Science and Technology (House of Representatives) was held in conjunction with a National Supercomputer Conference at the Florida State University in Tallahassee, Florida.

At this time, the Department of Energy sponsored supercomputer center at Florida State University had just gotten into production, and the National Science Foundation (NSF) sponsored centers were in various stages of doing so. Several of these NSF centers, along with private industry and federal and state governments, were represented at this conference.

This gathering confirmed that the major challenge that would be encountered in making the potential of supercomputer technology available to researchers would be a networking capability. Senator Albert Gore of Tennessee, speaking at the hearing, compared the need for a national networking capability to the need for the interstate highway system.

Subsequent development of the NSF plans for linking their supercomputer centers, along with discussions at conferences and the growth of local, state, and regional networks to serve both research and non-research needs, has resulted in a proliferation of networks that might be compared to the early proliferation of railroad networks after the invention of the steam railway engine. If we can carry that analogy a bit further, one problem in this early proliferation of railroad networks was the lack of a standard. Because of that, there are many areas of the world that even today experience a problem because of differing track sizes.

During the 99th Congress, 2nd Session, Senators Gore and Gorton introduced a bill that required the Office of Science and Technology Policy to report to the Congress on fiber optic networks and other options to improve communications among supercomputer centers and users in the United States.

## 21.2. PRESENT NEEDS

In considering general research needs, as opposed to specialized agency requirements, the NSF has already begun a promising program of encouraging network development with its program announcement for connection to NSF's National Supercomputer Access Network (NSFnet). The philosophy and development of NSFnet was set forth in an article in *Science* (Volume 231, February 28, 1986) by Dennis Jennings and others. Two extremely important results have emerged from this effort:

- NSF has enforced a standard protocol: TCP/IP.
- NSF has strongly encouraged the development and integration of local and regional nets into NSFnet.

In particular, the funding by NSF of SURAnet, a 56-kbit/s network connecting a consortium of universities throughout the southeastern United States, has been an important step. The philosophy of the group within SURA, which formulated this proposal (in which Florida State University participated), was that networking needs were much broader than merely supercomputer access, in fact much broader than computer access itself. This group, chaired by Dr. Jesse E. Poore, Jr., of Georgia Tech, theorized that synergism among researchers, when linked via a high-speed network, would result in total research productivity that would be far greater than the sum of the parts.

It must be kept in mind also that networking with collaborators overseas (i.e., Europe and Japan) is also very important, and may become even more so in the future.

## 21.3. FUTURE NEEDS AND THE ROLE OF FEDERAL GOVERNMENT

The subjects we should now focus our attention on, given this promising start, are a very high speed National Science Network for researchers, and the proper role that the federal government should play in its development. This role should be in three forms, described below.

- The Federal Government Should Fund and Install a High-Bandwidth, Redundant, Backbone Network Capability for Research Use.

Only the federal government has the resources to fund and install the type of national high-bandwidth backbone network that will be required to meet the needs of researchers. Although we cannot accurately forecast what the traffic will ultimately be on such a network, it would probably be well to remember that in the early 1960s Seymour Cray thought that two or three 6600 supercomputers would meet all the needs of researchers for computing. Such a backbone network should be redundant, not only to provide integrity of the system, but to provide extra capacity in case of unexpected growth.

- The Federal Government Should, Through Established Contract and Grant Mechanisms, Make it Attractive for Research-Oriented Universities to Interface with this Backbone Network.

It is not to be expected that researchers will, the minute such a network is available, immediately make use of the network. Yet, the thesis is that productive research within the United States will grow exponentially by the synergism among those researchers on the network. That being the case, the established contract and grant mechanisms of federal agencies should be such that it becomes attractive for a researcher to utilize the facilities. Such mechanisms vary from encouraging the sharing of instrumentation resources via the network, to encouraging the collaboration of geographically separated scientists on common research projects utilizing the network as a communications medium, to a requirement for the use of electronic mail via the



network for reporting purposes. The point is that there are many ways to encourage researchers to utilize the network. Once encouraged to do so, there is every reason to believe that researchers will demand access.

The model of a library (or interstate highways) should be used in funding the university connection to the network. It should be available to all (on a first-come, first-served basis) with the funding coming out of the overhead (or taxes). This would make it available to all on a continuous basis for scholarly pursuits without the need for unit charging (like a telephone company). These comments are suggested to direct that there not be an OMB Circular 21 prime for networking.

- The Federal Government Should Quickly Move to Establish an Accepted Protocol for Networks, and Make it Attractive for Regions and States to Utilize Such Protocol.

It is in the best long-term interests of the federal government not only to develop a high-speed backbone network as discussed above, but also to encourage the individual states and regions to have the capability to interface with the network. However, the development of such state and regional networks is driven by forces other than research needs. Because of varying needs, primary dependency upon certain types of computer equipment by state and regional networks does not necessarily favor the use of the same protocols and interface mechanisms. While it is not the scope of this paper to make recommendations as to the desirability of one protocol over another, it is desirable that there be a standard. The federal government is in the best position to encourage the development of such a standard.

The NSF program for connecting state and regional networks to NSFnet has already shown that there is a strong incentive for those states and regions that are driven by a commitment to research needs to conform to the standards required by NSFnet. This approach can be successfully utilized for the backbone network, if the incentives are sufficiently attractive.

The U.S. should choose a protocol that will serve the U.S. for the time being, but bear in mind that when an international standard is ready, we expect people to convert to it. Compatibility and interoperability with foreign (European and Japanese in particular) network protocols based on agreed-to international standards should be a stated goal, even though this may require another round of conversions.

In any event, the case for an accepted, common protocol is so strong as to completely override any opposing arguments. Only the federal government has the capability of implementing this common protocol and enforcing its usage by agencies.



## APPENDIX A -- ADDITIONAL REFERENCES

---

"An Agenda for Improved Evaluation of Supercomputer Performance," Report of the Committee on Supercomputer Performance and Development, National Research Council (1986).

Presents the results of a study to assess current methods of evaluating supercomputer performance, evaluate opportunities, and recommend research agenda. Recommends more support from funding agencies and leadership roles by NSF and DOE.

Proceedings of Workshop on Supercomputing Environments, June 24-26, 1986, NASA-Ames Research Center.

Includes papers on computer architecture, interactive and graphics environments, design and implementation of distributed storage systems, high-performance networks, satellite communications, and networking strategies and architecture.

D. M. Jennings, L. H. Landwebber, R. H. Fuchs, D. J. Farber, and W. R. Adrion, "Computer Networking for Scientists," *Science* **231**, (February 28, 1986).

Discusses scientists' needs for immediate access to data and information, to colleagues and collaborators, and to advanced computing and information services.

E. I. Holstrom, "Access to Supercomputers," Higher Education Panel Report Number 69, American Council on Education (January 1986).

A study sponsored by the NSF surveying 1,190 departments in all major research universities on computer use and assistance needed to increase access to supercomputers.

J. S. Quarterman and J. C. Hoskins, "Notable Computer Networks," *Communications of the ACM* **29**, 10 (October 1986).

Discusses the various characteristics of research networks, company networks, cooperative networks, metanetworks, and bulletin boards, and reviews legal and social issues.

Interface 88 Papers, Proceedings of Conference Co-sponsored by Week and Data Communications Magazines, March 4-7, 1985.

Includes papers on basic networking, data terminals, local networks, micro-to-mainframe issues, network design, network management and planning, regulation, security, standards, switching technology, testing and diagnostics, and voice and data networking.

J. Barley, "Personal Computer Networks," NBS Special Publication 500-140 (July 1986).

Surveys personal computer network technology from the point of view of the end user, and characterizes the capabilities of personal computer networks and services that they provide to the user in terms of generic features.

D. G. Perry, "Network Support of Supercomputers," Los Alamos National Laboratory conference report, June 5-7, 1985.

Summarizes discussion on existing supercomputer centers and services necessary to support users accessing supercomputer centers.

Proceedings for Computer Networking Plans, sponsored by Yale Computer Center; Proceedings for PC/IP MAC/IP Workshop, sponsored by the University of Maryland Computer Science Center, December 1985.

Includes copies of transparencies and a few of the papers delivered at the joint workshop. Subjects covered are plans for six universities to link together different host computers and workstations and activities at eleven universities to use microcomputers to link to the NSF internet.

D. J. Farber, G. Delp, M. Minnich, R. Minnich, G. Parulkar, and P. Schragger, "Research Directions in Networking Computers at the University of Delaware," University of Delaware Department of Electrical Engineering.

Discusses research projects focusing on future implications of advances in transmission and switching technology being developed at the communications industries' laboratories.

G. Delp and M. Minnich, "Some Future Directions at Bell Communications Research," University of Delaware, Udel-EE Technical Report Number 86-11-2 (November 7, 1986).

Reports on a meeting at Bellcore on October 24, 1986, attended by NSF, University of Delaware, and Bellcore staff. Topics discussed related to projected capabilities of national communications systems in the next decade.

Supercomputers: Government Plans and Policies. report of the Office of Technology Assessment, U.S. Congress. March 1986.

Presents a review of the Federal Government's large-scale computing programs and examines the networking and software programs within selected agencies. Certain management and institutional questions pertinent to the Federal efforts are also raised.

## APPENDIX B -- GLOSSARY

---

Definitions for the terms in this glossary (other than acronyms) are from the *Computer Dictionary*, Sipple and Sipple, Third Edition, and from *A GLOSSARY for Users of the Integrated Computing Network*, Los Alamos National Laboratory.

### A

ACI	The Advanced Computing Initiative graphics project at Los Alamos.
AI	Acronym for artificial intelligence.
ANSI	American National Standards Institute.
ARPANET	A resource-sharing network sponsored by the Advanced Research Projects Agency.
ASCII	American Standard Code for Information Interchange - an 8 bit code.

### B

BITNET	A resource sharing network sponsored by the participating sites.
bit/s	Bits per second.
broadband	Also called wideband. Pertaining to a channel with a bandwidth greater than voice-grade channels.
BSD	Berkeley Scientific Development
buses	Circuits over which data or power are transmitted.
bytes	Indicates a measurable portion of consecutive binary digits (bits). Common sizes are 6, 8, and 9 bits.
byte/s	Bytes per second.

### C

CCF	Central Computing Facility (Los Alamos).
CCITT	Consultative Committee for International Telephony and Telegraphy.
CC&R	Commitment, concurrency, and recovery.
CFD	Acronym for computational fluid dynamics.
CGM	Computer Graphics Metafile.
CPU	Acronym for central processor unit.
CSLAN	The Computing Services segment of the Lawrence Berkeley Laboratory network.
CSMA	Carrier Sense Multiple Access.

CSNET	A computer science researcher's network that provides a gateway to ARPANET/MILNET.
CSRLAN	The Computer Science Research segment of the Lawrence Berkeley Laboratory network.
CTSS	Cray Timesharing System.

**D**

DACS	Digital Automatic Cross-Connect Systems.
DARPA	Defense Advanced Research Projects Agency.
DCA	Defense Communications Agency.
DDS	Digital Data Service.
DEC	Digital Equipment Corporation.
DECnet	A resource-sharing network of DEC.
DES	Data Encryption Standards
dialup	The service whereby a dial telephone can be used to initiate and effect station-to-station telephone calls.
DOE	Department of Energy.
DoD	Department of Defense.
DP	Acronym for distributed processor.
DR11W	A communications interface component manufactured by DEC.
DS-0, DS-1, DS-3	Data transmission designations representing 64 kbit/s, 1.5 Mbit/s, and 45 Mbit/s data rates.
DSUWG	Data Systems Users Working Group.

**E**

EBDIC code	An eight-level code similar to the American Standard Code for Information Interchange. Abbreviation for expanded binary coded decimal interchange code.
EGP	Exterior Gateway Protocol.
EIN	European Information Network
ESnet	Energy Science Network of the DOE.
Ethernet	A high-speed local area networking technology based in IEEE Standard 802.3

**F**

<b>FCCSET</b>	Federal Coordinating Council on Science, Engineering and Technology.
<b>FDDI</b>	Fiber Distribution Data Interface.
<b>firmware</b>	Computer programs that are embodied in a physical device that can form part of a machine.
<b>FTP</b>	Acronym for file transfer protocol.

**G**

<b>GaAs</b>	Gallium arsenide.
<b>Gbit</b>	Abbreviation for gigabit, which is one billion bits.
<b>Gbyte</b>	Abbreviation for gigabytes.
<b>GGP</b>	Gateway-to-Gateway Protocol.
<b>GKS</b>	Graphical Kernel System.
<b>GSFC</b>	Goddard Space Flight Center.

**H**

<b>HDLC</b>	High-Level Data Link Control.
<b>HEPnet/LEP3NET</b>	Networks that Support High Energy Physics programs of DOE and NSF.
<b>HYPERchannel</b>	A registered trademark of Network Systems Corporation. Defines a high-speed data transmission bus.

**I**

<b>IBM</b>	International Business Machines.
<b>IEEE</b>	Institute of Electrical and Electronic Engineers.
<b>IFIP</b>	International Federation for Information Processing.
<b>IGP</b>	Interior Gateway Protocol.
<b>InteCom</b>	A digital switch from DEC.
<b>IP</b>	Internet Protocol.
<b>IRAB</b>	Internet Research Activities Board.
<b>IRI</b>	Interagency Research Internet.
<b>IRIO</b>	Interagency Research Internet Organization.
<b>IRIS</b>	Integrated Raster Imaging System workstation from Silicon Graphics Inc.

ISDN Integrated Services Digital Network.  
ISO International Standards Organization.  
ISP Gould Image Support Processor.

**J**

JPL Jet Propulsion Laboratory.

**K**

kbit Abbreviation for kilobit. A kilobit is one thousand binary bits.  
kbit/s Abbreviation for kilobits per second.  
kbyte Abbreviation for kilobytes.  
kHz Abbreviation for kilohertz.

**L**

LANs Local area networks.  
LaRC NASA's Langley Research Center.  
LeRC NASA's Lewis Research Center.  
LU 6.2 A general-purpose interprogram synchronous protocol for distributed applications developed by IBM.

**M**

mainframe The fundamental portion of a computer, i.e., the portion that contains the CPU and control elements of a computer system.  
MAN Metropolitan Area Network  
MAP Manufacturing Automation Protocol  
Mbit Abbreviation for million bits. An Mbit is one million binary bits.  
Mbit/s Abbreviation for million bits per second.  
Mbytes Abbreviation for million bytes per second.  
MFENET A communication network that supports the National Magnetic Fusion Energy Program.  
MFLOPS Million floating-point operations per second.  
MHz Abbreviation for megahertz.  
MILNET A resource-sharing network formerly a portion of ARPANET. The DoD Military Network.



MINET	A resource-sharing European network sponsored by DARPA.
MIPS	Million instructions per second.
MOA	Memo of authorization.
MSC	Mission Support Communications.
ms	Abbreviation for millisecond.
MSFC	Marshall Space Flight Center.
MVS	Memory Virtual System, a trademark of IBM.

## N

NAS	Numerical Aerodynamic Simulation (a NASA program).
NASA	National Aeronautics and Space Administration.
NASnet	A resource-sharing network sponsored by NASA.
NCS	National Communications System
NEC	Nippon Electron Corporation.
NESDIS	National Environmental Satellite and Data Information Services.
NeWS	Network Extensible Window System, trademark of Sun Microsystems, Inc.
NFS	Network File System.
NJE	IBM software that provides file transfer and electronic mail.
NMFECC	National Magnetic Fusion Energy Computing Center.
NOC	Acronym for network operation center.
NSF	National Science Foundation.
NSFnet	A resource-sharing network sponsored by NSF.
NSI	NASA Science Internet.
NTSC	National Television Standards Code.

## O

ODA	Office of Document Architecture.
ODC	Acronym for other direct costs.
OER	Office of Energy Research.
online	Descriptive of a system and peripheral equipment or devices in a system in which the operation of such equipment is under control of the CPU.
Op code	A command, usually given in machine language.

OSI Open Systems Interconnect.  
OSTP Office of Science and Technology Policy.

**P**

PBX Private branch exchange.  
PC Acronym for personal computers.  
PDP11 A minicomputer manufactured by DEC.  
PIN Acronym for personal identification numbers.  
PLOT3D Software used in plotting three-dimensional graphics.  
ProNET A high-speed network implemented at the University of Illinois.  
Proteon Proteon, Inc.  
PSC Program Support Communications.  
PSCN NASA's Program Support Communication Network.  
RGB Acronym for color television transmission of red, green, blue.  
RIACS Research Institute for Advanced Computer Science.  
RJE Remote Job Entry.  
RPC Remote Procedure Call.

**S**

SDLC Acronym for synchronous data link control.  
SDSC San Diego Supercomputer Center.  
SMTP A mail protocol.  
SNA System Network Architecture.  
SPAN Space Physics Analysis Network.  
SURAnet A network connecting a consortium of universities throughout the southeastern U.S.

**T**

T1 See DS-0, DS-1, DS-3.  
TAB Technical advisory board.  
TCP Transmission Control Protocol.  
Telnet Terminal access network.  
TEXnet Texas Academic Network.

Trans Transaction Network Service

## U

UCAR University Corporation for Atmospheric Research.  
UNICOS A UNIX operating system for the Cray computers.  
UNIX A timesharing operating system developed by Bell Laboratories.  
USENET A world-wide distribution network for electronic mail.  
USG Ultra-High-Speed Graphics program at Los Alamos.  
UUCP UNIX to UNIX Communication Protocol.

## V

VAX A family of mini-computers developed by Digital Equipment Corporation.  
VCR Video Cassette Recorder.  
VLSI Acronym for very large scale integration.  
VME Virtual Memory Extension, a new operating system and a trademark of IBM.  
VMS Virtual Memory System, a trademark of DEC.

## W

WDM Wavelength Division Multiplexing.

## X

XNS Xerox Network System.



---

**A RESEARCH AND DEVELOPMENT  
STRATEGY  
FOR  
HIGH PERFORMANCE COMPUTING**

---

---

Executive Office of the President  
Office of Science and Technology Policy  
November 20, 1987

EXECUTIVE OFFICE OF THE PRESIDENT  
OFFICE OF SCIENCE AND TECHNOLOGY POLICY  
WASHINGTON, D.C. 20506

This year the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) Committee on Computer Research and Applications began a systematic review of the status and directions of high performance computing and its relationship to federal research and development. The Committee held a series of workshops involving hundreds of computer scientists and technologists from academia, industry, and government. A result of this effort is the report that follows, containing findings and recommendations concerning this critical issue. It has been sent to the appropriate committees of Congress for their review.

A consistent theme in this report is the need for industry, academia, and government to collaborate and exchange information on future R&D efforts. Partners need to give one another signals as to their intent for future activities, and this report is a necessary first step in that process. The vision it represents must continue to grow. For that reason, I have asked the Committee to initiate the appropriate forums for discussing it further with the computing community.

Another theme has come out of this report: within four decades, the field of computer science has moved from a service discipline to a pervasive technology with a rigorous scientific basis. Computer science has become important to our national security and to our industrial productivity, and as such it provides the United States with many opportunities and challenges. Three of those opportunities are addressed in the report's findings and recommendations: High Performance Computers, Software Technology and Algorithms, and Networking. The fourth recommendation involves the Basic Research and Human Resources that will be required to conduct the other initiatives.

One thing is clear: the competition in an increasingly competitive global market cannot be ignored. The portion of our balance of trade supported by our high performance computing capability is becoming more important to the nation. In short, the United States must continue to have a strong, competitive supercomputing capability if it is to remain at the forefront of advanced technology. For that reason the Office of Science and Technology Policy is encouraging activities among the federal agencies together with the academic community and the private sector.



William R. Graham  
Science Adviser to the President and  
Director, Office of Science and Technology Policy

# CONTENTS

SUMMARY OF FINDINGS	1
SUMMARY OF RECOMMENDATIONS	2
THE CHALLENGE	3
THE STRATEGY	4
CURRENT STATUS AND TRENDS	5
IMPACT	7
BACKGROUND	8
1. HIGH PERFORMANCE COMPUTERS	12
2. SOFTWARE TECHNOLOGY AND ALGORITHMS	15
3. NETWORKING	18
4. BASIC RESEARCH AND HUMAN RESOURCES	23
IMPLEMENTATION	25
COST ESTIMATES	26
FCCSET COMMITTEE MEMBERS	29

# SUMMARY OF FINDINGS ON COMPUTER RESEARCH AND APPLICATIONS

- 1. HIGH PERFORMANCE COMPUTERS: A strong domestic high performance computer industry is essential for maintaining U.S. leadership in critical national security areas and in broad sectors of the civilian economy.**
  - U.S. high performance computer industry leadership is challenged by government supported research and development in Japan and Europe.
  - U.S. leadership in developing new component technology and applying large scale parallel architectures are key ingredients for maintaining high performance computing leadership. The first generation of scalable parallel systems is now commercially available from U.S. vendors. Application-specific integrated circuits have become less expensive and more readily available and are beginning to be integrated into high performance computers.
  
- 2. SOFTWARE TECHNOLOGY AND ALGORITHMS: Research progress and technology transfer in software and applications must keep pace with advances in computing architecture and microelectronics.**
  - Progress in software and algorithms is required to more fully exploit the opportunity offered by parallel systems.
  - Computational methods have emerged as indispensable and enabling tools for a diverse spectrum of science, engineering, design, and research applications.
  - Interdisciplinary research is required to develop and maintain a base of applications software that exploits advances in high performance computing and algorithm design in order to address the "grand challenges" of science and engineering.
  
- 3. NETWORKING: The U.S. faces serious challenges in networking technology which could become a barrier to the advance and use of computing technology in science and engineering.**
  - Current network technology does not adequately support scientific collaboration or access to unique scientific resources. At this time, U.S. commercial and government sponsored networks are not coordinated, do not have sufficient capacity, do not interoperate effectively, and do not ensure privacy.
  - Europe and Japan are aggressively moving ahead of the U.S. in a variety of networking areas with the support of concentrated government and industry research and implementation programs.
  
- 4. BASIC RESEARCH AND HUMAN RESOURCES: Federal research and development funding has established laboratories in universities, industry, and government which have become the major sources of innovation in the development and use of computing technology.**



# **SUMMARY OF RECOMMENDATIONS FOR A NATIONAL HIGH PERFORMANCE COMPUTING STRATEGY**

- 1. HIGH PERFORMANCE COMPUTERS:** The U.S. Government should establish a long range strategy for Federal support for basic research on high performance computer technology and the appropriate transfer of research and technology to U.S. industry.
- 2. SOFTWARE TECHNOLOGY AND ALGORITHMS:** The U.S. should take the lead in encouraging joint research with government, industry, and university participation to improve basic tools, languages, algorithms, and associated theory for the scientific “grand challenges” with widespread applicability.
- 3. NETWORKING:** U.S. government, industry, and universities should coordinate research and development for a research network to provide a distributed computing capability that links the government, industry, and higher education communities.
- 4. BASIC RESEARCH AND HUMAN RESOURCES:** Long term support for basic research in computer science should be increased within available resources. Industry, universities, and government should work together to improve the training and utilization of personnel to expand the base of research and development in computational science and technology.

---

# A RESEARCH AND DEVELOPMENT STRATEGY FOR HIGH PERFORMANCE COMPUTING

---

**High performance computing** refers to the full range of supercomputing activities including existing supercomputer systems, special purpose and experimental systems, and the new generation of large scale parallel architectures.

## THE CHALLENGE

In the span of four decades, computing has become one of the most pervasive and powerful technologies for information management, communications, design, manufacturing, and scientific progress.

The U.S. currently leads the world in the development and use of high performance computing for national security, industrial productivity, and science and engineering, but that lead is being challenged. Through an increased foreign industrial capability, the U.S. technology lead in computing has diminished considerably in recent years, but the U.S. continues to maintain strength in basic science and technology. The technology is changing rapidly and the downstream rewards for leadership are great. Progress in computing can be accelerated through the continued pioneering of new hardware, software, algorithms, and network technology and the effective transition of that technology to the marketplace. A shared computing research and development vision is needed to provide to government, industry, and academia a basis for cooperative action. The successful implementation of a strategy to attain this vision and a balanced plan for transition from one generation of technology to the next can result in continued strength and leadership in the forthcoming decades.

High performance computing technology has also become essential to progress in science and engineering. A **grand challenge** is a fundamental problem in science or engineering, with broad applications, whose solution would be enabled by the application of the high performance computing resources that could become available in the near future. Examples of grand challenges are: (1) Computational fluid dynamics for the design of hypersonic aircraft, efficient automobile bodies, and

extremely quiet submarines, for weather forecasting for short and long term effects, efficient recovery of oil, and for many other applications; (2) Electronic structure calculations for the design of new materials such as chemical catalysts, immunological agents, and superconductors; (3) Plasma dynamics for fusion energy technology and for safe and efficient military technology; (4) Calculations to understand the fundamental nature of matter, including quantum chromodynamics and condensed matter theory; (5) Symbolic computations including speech recognition, computer vision, natural language understanding, automated reasoning, and tools for design, manufacturing, and simulation of complex systems. Many of these could be considerably advanced by the use of computer systems capable of trillions of operations per second.

## **THE STRATEGY**

**A High Performance Computing Strategy**, involving close coordination of existing programs and augmented effort, is required to address this national challenge. This strategy involves the coordinated pursuit of computing technology goals through joint efforts of government, industry, and academia. The strategy will have impact in clarifying and focusing the direction of Federally-funded computing research, which continues to be the major source of innovation for computing technology and a primary catalyst for industrial development. Government support should be highly leveraged with resources provided by industry participants. To be effective, the strategy should also be defined and continually updated in cooperation with industry and academia by making them participants in developing and implementing a shared vision of the future to ensure continued U.S. leadership.

The high performance computing strategy is designed to sustain and focus basic Federally-funded research and promote the transfer of basic science from the laboratory to U.S. industrial development and finally to the marketplace. Technology development will be encouraged as appropriate to meet immediate needs as well as to create a foundation for long term leadership. Strong emphasis will be placed on continued transfer of the results of government funded R&D to industry and on cooperation with industry to insure the continued strength of American high technology trade in the international marketplace.

The basic elements of the strategy are research and development programs in high performance computer architecture, in custom hardware, in software and algorithms, and in networking technology, all supported by a basic research foundation. In each of these areas, major opportunities exist that require coordinated support and management, building on existing government programs. Access to high performance computing is essential for providing scientists and engineers at research institutions throughout the country with the ability to use the most advanced computers for their work. The strategy needs to concurrently address the appropriate Federal role in each

of the basic elements of the R&D process—basic research, applied research, and industrial development—in order to meet long term, intermediate, and short term technology development goals. Explicit attention must be directed to the flow of technology from basic to applied areas and to the marketplace, as well as back into the research community to create the next generation of computing infrastructure, achieving a cumulative effect. Technology developments within individual element areas will contribute extensively to other activities. Simultaneous and coordinated pursuit of the areas is therefore an important element of the strategy.

## CURRENT STATUS AND TRENDS

- **High performance computing systems.** Improvements in materials and component technology are rapidly advancing computer capability. Memory and logic circuits are continuing to improve in speed and density, but as fundamental physical limits are approached, advances are being sought through improved computer architectures, custom hardware, and software. Computer architecture has begun to evolve into large scale multiple processor systems, and in the past four years a first generation of scalable parallel systems has progressed from the research laboratory to the marketplace. Scalable architectures provide a uniform approach that enables a wide range of capacity, from workstations to very high performance computers. Application-specific integrated circuits, such as for real-time signal processing, are being incorporated into special purpose computers.

At current performance levels our ability to model many important science, engineering, and economic problems is still limited. Formulations of computational models presently exist that for realistic solutions would require speeds of teraflops (trillions of floating point operations per second) and equivalent improvement in memory size, mass storage, and input/output systems. In addition, symbolic processing is complementing and enhancing numeric approaches. Achievement of this performance level in the next 5 years appears to be a feasible goal, based on credible extrapolations of processor capability, number of processors, and software sophistication. In developing the new architectural approaches, however, careful collaboration will be required with the applications community to assess the various approaches and to achieve transition to the new approaches where appropriate. As transitions are made, the high performance computing industry should strive to maintain its continued leadership and competitiveness.

- **Software technology and algorithms.** As high performance computing systems evolve and become more critical in science, engineering, and other applications domains, software technology becomes an increasingly central concern. As experienced in many U.S. space and defense programs, for example, software can become the dominant computational cost element in large systems because of the need to support evolution throughout the system life cycle from design and

development to long term maintenance and transition to the next generation. Future software environments and tools should support the development of trustworthy systems capable of evolution, while increasing productivity of developers and users of the systems. Effective exploitation of the performance potential of the emerging parallel systems poses a special challenge both to software and to algorithm design.

High performance computing offers scientists and engineers the opportunity to use computer models to simulate conditions difficult or impossible to create and measure in the laboratory. This new paradigm of computational science and engineering offers an important complement to traditional theoretical and experimental approaches, and it is already having major impact in many areas. New approaches combining numeric and symbolic methods are emerging. The development of new instruments and data generation methods in fields as diverse as genetics, seismology, and materials accelerates demand for computational power. In addition, the opportunity is created to coordinate and focus effort on important grand challenges, such as computational fluid dynamics, weather forecasting, plasma dynamics, and other areas.

- **Computer network technology.** A modern high speed research network is one of the elements needed to provide high performance distributed computation and communication support for research and technology development in government, academia, and industry. A coordinated research network based on very high bandwidth links would enable the creation of large-scale geographically distributed heterogeneous systems that link multiple high performance workstations, databases, data generation sources, and extremely high performance servers as required, in order to provide rapid and responsive service to scientists and engineers distributed across the country. The existing national network is a collection of loosely coupled networks, called an internet, based on concepts pioneered in the U.S.

Technical issues being addressed include utilization of fiber optics to improve performance for the entire research and higher education enterprise of the nation. An additional issue of pressing concern, particularly within the governmental and industrial sectors, is that of computer and network security to ensure privacy and trustworthiness in a heterogeneous network environment. At present, responsibility for privacy and the assurance of trust are vested principally in the computers and switching nodes on the network. Further research, already actively underway, is urgently needed to develop models, methodology, algorithms and software appropriate to the scale of a coordinated research network.

- **Basic research and human resources in Computer and Computational Science.** Federal funding has historically been, and will likely remain, a major source of support for important new ideas in computing technology. Carefully managed and stable funding is required to maintain vigorous research in computer and computational science and sufficient growth in computer science manpower. It is important to maintain the strength of the existing major research centers and to develop new research activity to support the growth in computer and computational science. Interactions should be fostered among academia, industry, and national laboratories to address large problems and to promote transfer of technology. In the longer term, enhancement of the computing technology base will have significant impact in productivity, efficiency, and effectiveness of government, industry, and the research community.

## **IMPACT**

Computing technology is vital to national security. Advanced computer systems and software are now integral components in most major defense, intelligence, and aerospace systems. Computing technology has a central role in energy research, oil exploration, weapons research, aircraft design, and other national security technology areas.

Major advances in science and engineering have also accrued from recent improvements in supercomputing capability. The existence of machines with hundred megaflop (hundreds of millions of floating point operations per second) speed and multimillion word memories has allowed, for the first time, accurate treatment of important problems in weather prediction, hydrodynamics, plasma physics, stress analysis, atomic and molecular structure, and other areas. The emerging machines with 1 to 10 gigaflop (billions of flops) speed and 100 to 300 million word memories are expected to produce comparable advances in solving numeric and symbolic problems.

Many of these advances in science and engineering are the result of the application of high performance computing to execute computational simulations based on mathematical models. This approach to science and engineering is becoming an important addition to traditional experimental and theoretical approaches. In applications such as the National Aerospace Plane, supercomputing provides the best means to analyze and develop strategies to overcome technical obstacles that determine whether the hypersonic vehicle can fly beyond speeds of Mach seven, where wind tunnels reach their maximum capability. The list of applications for which supercomputing plays this kind of role is extensive, and includes nearly all high-technology industries. The extent of its usage makes supercomputing an important element in maintaining national competitiveness in many high technology industries.

The high performance computing strategy will have impact in many sectors of the economy. Nearly all sectors of advanced industry are dependent on computing infrastructure. Any improvement in computing capability will have substantial leveraged impact in broad sectors, particularly as applications software increases in power and sophistication.

The computer hardware industry alone amounted to \$65 billion in 1986, and U.S. technical market dominance, long taken for granted, is now challenged in this and other critical areas, including networking, microsystems and custom high-performance integrated circuit technology. Foreign investment in computing research and technology has grown considerably in the last decade.

As stated in the report of the White House Science Council, *Research in Very High Performance Computing*, November 1985, "The bottom line is that any country which seeks to control its future must effectively exploit high performance computing. A country which aspires to military leadership must dominate, if not control, high performance computing. A country seeking economic strength in the information age must lead in the development and application of high performance computing in industry and research."

## **BACKGROUND**

The Federal Coordinating Council on Science, Engineering and Technology (FCCSET) was established by Congress under the Office of Science and Technology Policy (OSTP) to catalyze interagency consideration of broad national issues and to coordinate various programs of the Federal government. The FCCSET in turn, established a series of committees, with interagency participation to assess and recommend action for national science and technology issues. The committees have become recognized as focal points for interagency coordination activity, addressing issues that have been identified by direct requests through the OSTP and indirect requests by member agencies (such as the NSF requirement to provide an update to the Lax Report on Large Scale Computing in Science and Engineering). These studies have enabled the FCCSET Committee on Computer Research and Applications to develop a national view of computing technology needs, opportunities, and trends.

From its inception, the FCCSET Committee on Supercomputing (the original name of this committee) was chartered to examine the status of high performance computing in the U.S. and to recommend what role the Federal Government should play regarding this technology. The committee issued two reports in 1983 that provided an integrated assessment of the status of the supercomputer industry and recommended government actions. The FCCSET Committee on Computer Research and Applications concluded that it would be proper to include an update of the earlier reports to address the changes that have occurred in the intervening period as a complement to the technical

reports. The review was based upon periodic meetings with and site visits to supercomputer manufacturers and consultation with experts in high performance scientific computing. White papers were contributed to this report by industry leaders and supercomputer experts. The report was completed in September 1987 and its findings and recommendations are incorporated in the body of this report.

In developing the recommendations presented in this report, the FCCSET Committee reviewed findings and recommendations from a variety of sources, including those mentioned above. A related activity has been the preparation by the White House Science Council (WHSC) Committee on Research in Very High Performance Computing of the report *Research in Very High Performance Computing*, November 1985. The WHSC Committee, composed of respected individuals from academia, industry, and government, made recommendations related to the issues more recently addressed by the FCCSET Committee. In the areas addressed by both committees, there is a significant consistency of recommendations, and, indeed, progress in recent months further strengthens the case for the recommendations. The convergence of views expressed in the many reports, the strong interest in many sectors of government in developing a policy, the dramatic increase in foreign investment and competitiveness in computing and network technology, and the considerable progress in computing technology development worldwide are all indicators of the urgency of developing and implementing a strategy for nationwide coordination of high performance computing under the auspices of the government.

One of the of the direct requests that this report responds to is in Public Law 99-383, August 21, 1986, in which Congress charged the Office of Science and Technology Policy to conduct a study of critical problems and of current and future options regarding communications networks for research computers, including supercomputers, at universities and federal research facilities in the United States. The legislation asked that requirements for supercomputers be addressed within one year and requirements for all research computers be addressed within two years. Dr. William R. Graham, Director of the Office of Science and Technology Policy, subsequently charged the Federal Coordinating Council on Science Engineering and Technology (FCCSET) Committee on Computer Research and Applications to carry out the technical aspects of the study for OSTP.

It was recognized by the FCCSET Committee on Computer Research and Applications that networking technology needs to be addressed in the context of the applications of computing and the sources of computing power that are interconnected using the network technology. This report, therefore, presents an integrated set of findings and recommendations related to Federal support for computer and related research.



Three subcommittees carried out the work. Each of these committees contributed to the Findings and Recommendations contained in this report. The result is an integrated set of recommendations that addresses the technical areas.

- The **Subcommittee on Computer Networking, Infrastructure, and Digital Communications** invited experts in government, industry and academia to write white papers on networking trends, requirements, concepts applications, and plans. A workshop involving nearly 100 researchers, network users, network suppliers, and policy officials was held in San Diego, California in February 1987 to discuss the white papers and to develop the foundation for the report. Workshop leaders and other experts later met in Washington to summarize the workshop discussions and focused on six topics: access requirements and future alternatives, special requirements for supercomputer networks, internet concepts, future standards and services requirements, security issues, and the government role in networking. As a result of this work, the participants recommended that no distinction should be made between networks for supercomputers and other research computers and that the final report to the Congress should address networks generally. The requirements for both supercomputers and for other research computers are, therefore, addressed in this report.

- The **Subcommittee on Science and Engineering Computing** assessed computing needs related to computational science and engineering. The committee focused its deliberations on requirements for high performance computing, on networking and access issues, and on software technology and algorithms. Under the auspices of the Society for Industrial and Applied Mathematics (SIAM), and with the support of NSF and DOE, a workshop involving 38 recognized leaders from industry, academia, and national laboratories was held at Leesburg, Virginia in February 1987 on research issues in large-scale computational science and engineering. This workshop focused on advanced systems, parallel computing and applications. As a result of the workshop report, recommendations were made related to the role of computing technology in science and engineering applications.

- The **Subcommittee on Computer Research and Development** assessed the role of basic research, the development of high performance computing technology, and issues related to software technology. Contributing to this activity were two workshops. The National Science Foundation (NSF) Advisory Committee for Computer Research reviewed the field and produced an Initiatives Report in May 1987. This report recommended investment in three areas, including parallel systems and software technology. In September 1987, the Defense Advanced Research Projects Agency (DARPA) held a workshop on advanced computing technology in Gaithersburg, Maryland involving 200 researchers from academia, industry, and government. The workshop focused on large-scale parallel systems and software approaches to achieving high performance computing.

An important result of the activity of the FCCSET Committee on Computer Research and Applications and its subcommittees is that increased coordination among the Government elements is necessary to implement a strategy for high performance computing. The findings and recommendations presented here represent a consensus reached among the subcommittees and convey the powerful and compelling vision that emerged. As a result of this process, the next step would be for the members of the Committee on Computer Research and Applications to develop a plan to help ensure that the vision is shared between government, academia, and American industry. Subsequently, the Committee should develop an implementation plan for Federal government activities, including a detailed discussion of overall priorities.

# 1. HIGH PERFORMANCE COMPUTERS

- **FINDING:** A strong domestic high performance computer industry is essential for maintaining U.S. leadership in critical national security areas and in broad sectors of the civilian economy.

U.S. prominence in technology critical to national defense and industrial competitiveness has been based on leadership in developing and exploiting high performance computers. This preeminence could be challenged by dependency upon other countries for state of the art computers. Supercomputer capability has contributed for many years to military superiority. In addition, industrial applications now constitute more than half of the supercomputer market and are an important factor in U.S. industrial competitiveness. However, continued progress in computational science and engineering will depend in large part on the development of computers with 100 to 1000 times current capability for important defense, scientific, and industrial applications. These applications are represented by the grand challenges.

- **U.S. high performance computer industry leadership is challenged by government supported research and development in Japan and Europe.**

The U.S. currently leads the world in research, development, and use of supercomputers. However, this leadership faces a formidable challenge from abroad, primarily from the Japanese. The 1983 FCCSET report stated that "The Japanese have begun a major effort to become the world leader in supercomputer technology, marketing, and applications." Most of the analyses and projections advanced in support of that statement have proven to be accurate.

Japanese supercomputers have entered the marketplace with better performance than expected. Japanese supercomputer manufacturers have attained a high level of excellence in high speed, high density logic and memory microcircuits required for advanced supercomputers. As a result, some U.S. computer manufacturers are dependent on their Japanese competitors for sole supply of critical microcircuits. Japanese manufacturers, universities, and government have demonstrated the ability to cooperate in developing and marketing supercomputers as well as in advancing high performance computing. Recent successes in dominating related high-technology markets underscore their financial, technical, and marketing capability.

- **U.S. leadership in developing new component technology and applying large scale parallel architectures are key ingredients for maintaining high performance computing leadership. The first generation of scalable parallel systems is now commercially available from U.S. vendors. Application-specific integrated circuits have become less expensive and more readily available and are beginning to be integrated into high performance computers.**

The current generation of supercomputers achieve their performance through the use of the fastest possible individual components, but with relatively conservative computer architectures. While these computers currently employ up to eight parallel processors, their specific architectures cannot be scaled up significantly. Large scale parallel processing, in which the computational workload is shared among many processors, is considered to be the most promising approach to producing significantly faster supercomputers. The U.S. is currently the leader in developing new technology as well as components. However, exploiting these techniques effectively presents significant challenges. Major effort will be required to develop parallel processing hardware, algorithms, and software to the point where it can be applied successfully to a broad spectrum of scientific and engineering problems.

Government funded R&D in universities and industry has focused on an approach to large-scale parallelism that is based on aggressive computer architecture designs and on high levels of circuit integration, albeit with somewhat slower individual components. Unlike current supercomputers, the resulting systems employ 100s to 10,000s of processors. Equally important, these architectures are scalable to higher levels of parallelism with corresponding increase in potential performance.

The first generation of scalable parallel systems is now commercially available from U.S. vendors. These systems have demonstrated high performance for both numeric and non-numeric, including symbolic processing. Comparable systems do not yet exist outside the U.S. The second generation, with higher speed individual components and more parallelism, is already in development here. Experience with these systems has shown that, even with existing software, they are effective for certain classes of problems. New approaches to software for these large-scale parallel systems are in the process of emerging. These approaches suggest that parallel architecture may be effective for wide classes of scientific and engineering problems. An important benefit of the scalable architectures is that a single design, with its attendant components and software, may prove to be useful and efficient over a performance range of 10 to 100 or more. This allows one design to be used for a family of workstations, mini-supercomputers, and supercomputers.

- **RECOMMENDATION:** The U.S. Government should establish a long range strategy for Federal support for basic research on high performance computer technology and the appropriate transfer of research and technology to U.S. industry.

The program should build upon existing government supported efforts. However, government funding should not be viewed as a substitute for private capital in the high performance computer marketplace. A primary objective is to ensure continued availability of domestic sources for high performance computers that are required for Federal programs, both civilian and defense. These actions should include:

- Government should support, when appropriate for mission requirements, the acquisition of prototype or early production models of new high performance computers that offer potential for improving research productivity in mission areas. These computers could be placed in centers of expertise in order to allow sophisticated users to share initial experiences with manufacturers and other users, and to develop software to complement the vendor's initial offerings. These initial acquisitions should not require the vendor to supply mature operating systems and applications software typical of production computers. However, a criterion for acquisition should be that the hardware designs reflect a sensitivity to software issues, and that the computer has the potential for sustained performance in practical applications that approaches the peak hardware performance.
- Government agencies should seek opportunities to cooperate with industry in jointly funded R&D projects, concentrating especially on those technologies that appear scalable to performance levels of trillions of operations per second (teraops) for complex science, engineering, and other problems of national importance. Systems are needed for both numeric and symbolic computations.

However, since government mission requirements typically exceed those of industrial applications, cooperating with industry in R&D for computers to meet these missions will help to assure that the necessary computers are available. It will also drive supercomputer development at a faster pace than would be sustained by commercial forces alone, an important factor retaining and increasing U.S. leadership in this area.

- Government agencies should fund basic research to lay the foundation for future generations of high performance computers. Steps should be taken to ensure that development of state-of-the-art computers continues to be monitored for appropriate export controls.

## 2. SOFTWARE TECHNOLOGY AND ALGORITHMS

- **FINDING:** Research progress and technology transfer in software and applications must keep pace with advances in computing architecture and microelectronics.
  - Progress in software and algorithms is required to more fully exploit the opportunity offered by parallel systems.
  - Computational methods have emerged as indispensable and enabling tools for a diverse spectrum of science, engineering, and design research and applications.
  - Interdisciplinary research is required to develop and maintain a base of applications software that exploits advances in high performance computing and algorithm design in order to address the “grand challenges” of science and engineering.

A **grand challenge** is a fundamental problem in science and engineering, with broad application, whose solution will be enabled by the application of the high performance computing resources that could become available in the near future.

As high performance computing systems evolve and are applied to more challenging problems, it is becoming increasingly clear that advances in software technology and applications are essential to realize the full performance potential of these systems. Software development, analysis, and adaptation remain difficult and costly for traditional sequential systems. Large scale complex systems including parallel systems pose even greater challenges. Market pressures for the early release of new computing system products have created a tradition of weak systems software and inadequate programming tools for new computers.

Current approaches to software development provide only limited capabilities for flexible, adaptable, and reusable systems that are capable of sustained and graceful growth. Most existing software is developed to satisfy nearer term needs for performance at the expense of these longer term needs. This is particularly the case for applications in which specific architectural features of computers have been used to obtain maximum performance through low level programming techniques. The lack of portability of these programs significantly raises the cost of transition to newer architectural approaches in many applications areas. Approaches are beginning to emerge in the research community that have a potential to address the reuse and portability problems.

Experiments with parallel computers have demonstrated that computation speeds can increase almost in direct proportion to the number of processors in certain applications. Although it is not yet possible to determine in general the most

efficient distribution of tasks among processors, important progress has nonetheless been made in the development of computational models and parallel algorithms for many key problem areas.

Access to advanced computing systems is an important element in addressing this problem. Experience has shown that the quality of systems and applications software increases rapidly as computing systems are made more available. Initial generic operating systems and extensions to existing programming languages can provide access through coupling high performance computers with existing workstations using either direct or network connections. However, in order to achieve the full potential impact of large scale parallel computing on applications, major new conceptual developments in algorithms and software are required.

The U.S. leads in many areas of software development. The Japanese, however, also recognize the need for high quality software capability and support in order to develop and market advanced machines. They have demonstrated the ability to effectively compete, for example in the area of sophisticated vectorizing compilers.

The U.S. will need to encourage the collaboration of computer scientists, mathematicians, and the scientists in critical areas of computing applications in order to bring to bear the proper mix of expertise on the software systems problem. Such collaboration will be enhanced by network technology, which will enable geographically dispersed groups of researchers to effectively collaborate on "grand challenges." Critical computer applications include problems in fluid dynamics, plasma physics, elucidation of atomic and molecular structure, weather prediction, engineering design and manufacturing, computer vision, speech understanding, automated reasoning, and a variety of national security problems.

- **RECOMMENDATION:** The U.S. should take the lead in encouraging joint research with government, industry, and university participation to improve basic tools, languages, algorithms, and associated theory for the scientific “grand challenges” with widespread applicability.

Software research should be initiated with specific focus on key scientific areas and on technology issues with widespread applicability. This research is intended to accelerate software and algorithm development for advanced architectures by increased early user access to prototype machines. It would also provide settings for developing advanced applications for production machines. Software technology needs to be developed in real problem contexts to facilitate the development of large complex and distributed systems and to enable transition of emerging parallel systems technology into the computing research community and into the scientific and engineering applications communities.

As part of a mixed strategy, longer term and more basic software problems of reliability and trust, adaptability, and programmer productivity must continue to be addressed. Languages and standards must be promoted that permit development of systems that are portable without sacrificing performance.

In applications areas including computational science and engineering, technology should be developed to support a smooth transition from the current software practice to new approaches based on more powerful languages, optimizing compilers, and tools supported by algorithm libraries. The potential of combining symbolic and numeric approaches should be explored. Progress in these areas will have significant impact on addressing the “grand challenges” in computational science and engineering. Although there are many pressing near term needs in software technology, direct investment in approaches with longer term impact must be sustained if there is to be significant progress on the major challenges for software technology while achieving adequate system performance.



### 3. NETWORKING

- **FINDING:** The U.S. faces serious challenges in networking technology which could become a barrier to the advance and use of computing technology in science and engineering.
  - Current network technology does not adequately support scientific collaboration or access to unique scientific resources. U.S. commercial and government sponsored networks presently are not coordinated, do not have sufficient capacity, do not interoperate effectively, and do not ensure privacy.
  - Europe and Japan are aggressively moving ahead of the U.S. in a variety of networking areas with the support of concentrated government and industry research and implementation programs.

Computer network technology provides the means to develop large scale distributed approaches to the collaborative solution of computational problems in science, engineering, and other applications areas. Today, researchers sharing a local area network are able to exploit nearly instantaneous communication and sharing of data, creating an effect of linking their workstations and high performance servers into a single large scale heterogeneous computing facility. This kind of capability is now appearing in larger scale campus-wide computer networks, enabling new forms of collaboration. National networks, on the other hand, have low capacity, are overloaded, and fail to interoperate successfully. These have been expanded to increase the number of users and connections but the performance of the underlying network technology has not kept pace with the increased demands. Therefore, the networks which in the 1970s had significant impact in enabling collaboration, are now barriers. Only the simplest capabilities, such as electronic mail and small file transfers, are now usable. Capacity, for example, is orders of magnitude less than the rates required, even if the network is used only for graphics.

Other countries have recognized the value of national computing networks, and, following the early U.S. lead, have developed and installed national networks using current technology. As a result, these countries are now much better prepared to exploit the new opportunities provided by distributed collaborative computing than the U.S. is at the present time. The basic technologies for later generations are also being developed in the U.S., but there have been no major efforts to apply them to address the needs.

Applications include (1) distributed access to very large databases of scientific, engineering, and other data, (2) high bandwidth access to and linking among shared computational resources, (3) high bandwidth access to shared data generation resources, (4) high bandwidth access to shared data analysis resources, such as workstations supporting advanced visualization techniques.

A longer term goal is the creation of large scale geographically distributed heterogeneous systems that link multiple superworkstations and high performance supercomputers to provide service to scientists and engineers distributed across the country. A well-coordinated national network could link these resources together when required on an *ad hoc* basis to provide rapid response to computational needs as they arise. This could reduce the number of sites needed for the physical presence of supercomputers. Present access to computer networks by researchers is dependent upon individual funding or location. There is unnecessary duplication in the links from various agencies to each campus. The development of improved networking facilities could greatly stimulate U.S. research and provide equitable access to resources.

Many scientific research facilities in the U.S. consist of a single, large, and costly installation such as a synchrotron light source, a supercomputer, a wind tunnel, a particle accelerator, or a unique database. These facilities provide the experimental apparatus for groups of scientific collaborators located throughout the country. Wide area networks are the logical mechanism for making data from such facilities more easily accessible nationwide. An important issue is that of computer and network security to ensure privacy and trustworthiness in a heterogeneous network environment. At present, responsibility for privacy and the assurance of trust are vested principally in the computers and switching nodes on the network.

Existing government-supported wide-area networks include ARPANET, HEPNET, MFENET, NSFNET, NASNET, MILNET, and SPAN, as well as private and commercial facilities such as TYMNET, TELENET, BITNET, and lines leased from the communication carriers. Longer-range estimates vary, but it is expected that by the year 1995 the nation's research community will be able to make effective use of a high capacity national network with capacity measured in billions of bits per second. Without improved networks, speed of data transmission will be a limiting factor in the ability of researchers to carry out complex analyses. The digital circuits most widely available today with transmission speeds of 56 kilobits per second (kb/s) are impediments to leading edge research and to optimal remote high performance computer use.

Point-to-point connections require interconnects through multiple vendors with cumulative costs. Greater network speed can reduce the time required to perform a given experiment and increase both the volume of data and the amount of detail that can be seen by researchers. Scientists accessing supercomputers would benefit because access speed is often critical in their work. Improved functionality frees scientists to concentrate directly on their experimental results rather than on operational details of the network. Increased network size extends these opportunities to thousands of individuals at smaller academic institutions throughout the nation. These modernization measures would significantly enhance the nation's position in scientific research. A national network would help maintain the U.S. leadership position in computer architectures,

microprocessors, data management, software engineering, and innovative networking facilities, and promote the development of international networking standards based on U.S. technology.

Integrated Systems Digital Networks (ISDN—voice and data) have been installed abroad on a national or regional scale. Research abroad is being conducted on service up to 1 Gb/s. Within the next five years, Integrated Services Digital Network (ISDN) circuits ranging from 64 kb/s to 1.5 Mb/s will be available in the larger metropolitan areas of the U.S. However, these services will fall short of the requirements for computer networks. By 1988 more than fifty Campus Area Networks will be operational at speeds approaching 100 Mb/s. Wide area networks operating at 1.5 Mb/s or less will not be able to handle the data volume expected.

Japan and Europe have extensive efforts with experimental nets in intermediate (40Mb) and high (gigabit) range. Japan is studying operational aspects of fiber nets using their national research network as a testbed, which includes exploring the feasibility of fiber optic services to residences.

To estimate the network bandwidth needed to support research at a major installation, the kinds and volume of traffic that would be used have been estimated at a representative campus, extrapolated ten years into the future. Three models were used to compute three independent estimates of the requirements for bandwidth needed by type of work, information needs by type of user, and information flowing at the installation boundary. In each model, the peak bandwidth was estimated for each type of service. For example, in the Task model, the need is dominated by that of at least one researcher to receive full color and full-motion high resolution images. A high-resolution color image contains about 30 megabits of information, so that a display rate of 30 frames per second requires a bandwidth of nearly one gigabit per second (Gb/s). In the User model, a research university with 35,000 students and 3,000 faculty and research staff using a mix of bandwidths again requires an aggregate bandwidth of approximately one Gb/s. In the Edge of the Installation model, bandwidth is estimated by the types of remote facilities being accessed and the expected number of simultaneous users; typical facilities include particle accelerators, supercomputers, and centers for imaging and/or animation. The aggregate bandwidth needed is one Gb/s. Thus three independent means of estimating bandwidth arrive at nearly the same requirement for a large research installation, and one Gb/s can confidently be used as a lower bound on the bandwidth of a national research network.

● **RECOMMENDATION: U.S. government, industry, and universities should coordinate research and development for a research network to provide a distributed computing capability that links the government, industry, and higher education communities.**

A research network should be established in a staged approach that supports the upgrade of current facilities and development of needed new capabilities. Achievement of this goal would foster and enhance the U.S. position of world leadership in computer networking as well as provide infrastructure for collaborative research. The FCCSET Committee on Computer Research and Applications should provide a forum for interagency cooperations. Elements of the plan should include:

- *Stage 1.* Upgrade existing facilities in support of a transition plan to the new network through a cooperative effort among major government users. The current interagency collaboration in expanding the Internet system originated by DARPA should be accelerated so that the networks supported by the agencies are interconnected over the next two years.
- *Stage 2.* The nation's existing networks that support scientific research should be upgraded and expanded to achieve data communications at 1.5 Mb/sec for 200 to 300 U.S. research institutions.
- *Stage 3.* Develop a system architecture for a national research network to support distributed collaborative computation through a strong program of research and development. A long-term program is needed to advance the technology of computer networking in order to achieve data communication and switching capabilities to support transmission of three billion bits per second (3 Gb/s) with deployment within fifteen years.
- Develop policy for long term support and upgrading of current high performance facilities, including timetables for backbone and connection development, industry participation, access, agency funding, tariff schedules, network management and administration. Support should be given to the development of standards and their harmonization in the international arena.

Until the national research network can replace the current system, existing networks should be maintained and modified as they join the national network. Remedial action should be initiated as soon as possible. Upgrading the backbone to at least 1.5 Mb/s should be accomplished by 1990. This will ensure that the new generation of high performance computing can be effectively interconnected.

Industry should be encouraged to participate in research, development, and deployment of the national research network. Telecommunication tariff schedules

which have been set for voice transmission should be reviewed in light of the requirements for transmission of data through computer networking.

Prompt effective coordination is needed to increase user participation in the standards development process, to get requirements for standards expressed early in the development process, and to speed the implementation of standards in commercial off-the-shelf products. It is essential that standards development be carried out within the framework of overall systems requirements to achieve interoperability, common user interfaces to systems, and enhanced security.

## 4. BASIC RESEARCH AND HUMAN RESOURCES

- **FINDING:** Federal research and development funding has established laboratories in universities, industry, and government which have become the major sources of innovation in the development and use of computing technology.

Many of the advances in computer science and technology in the U.S. were made possible by Federal programs of research support to universities and industry. For example, the advances that have occurred since 1983 in the area of parallel computing are the direct result of Federal research investment through agencies including DARPA and NSF. In the area of application of supercomputers to science and engineering, the majority of this investment came from the NSF Advanced Scientific Computing centers. In the area of parallel architectures, the major investment came from the DARPA Strategic Computing Program. Programs sponsored by DOE, NASA, and Defense to support critical mission needs have been a major source of investment in computational applications research. In industry, support for basic research is only a small fraction of industry research most of which is focused on nearer term product development. This can be attributed in part to the long term and high risk nature of basic research, but a more significant inhibitor of investment is the difficulty in the computer industry of maintaining proprietary protection for certain kinds of key fundamental advances.

- **RECOMMENDATION:** Long term support for basic research in computer science should be increased within available resources. Industry, universities, and government should work together to improve the training and utilization of personnel to expand the base of research and development in computational science and technology.

Maintain vigorous research in Computer Science and sufficient growth of computer science manpower to support the scientific/technological basis of the computer field. Foster interactions among academia, industry, and national laboratories by creating interdisciplinary teams to address large scale problems. Extend the technology base to attain significant impact on competitiveness and industrial productivity.

Innovative very high performance computing systems should be made available to universities and basic research laboratories in order to assist in the evaluation and exploitation of new technology and new industrial innovations.

Continue the following successful approaches to basic research and development: (1) The practice of loosely coordinated and flexible basic research supported through various federal sectors and applied to a diversity of institutions, (2) The mixed strategy of peer review to support a broad range of exploratory basic research throughout the academic community and the complementary technical program management approach of larger scale experimental systems programs which exploit new opportunities as they emerge, (3) Support for individuals and small groups in theoretical areas, (4) The practice of supporting the relevant basic research as part of larger experimental systems projects.

## IMPLEMENTATION

Success of the National High Performance Computing Strategy will require an attitude of cooperation in which academia, industry and government work effectively together in developing and assessing new technology and in achieving the transition of promising new ideas into the marketplace. The rapid pace of developments in computing technology creates a number of implementation challenges that must be addressed explicitly if the Strategy is to have maximum impact.

The FCCSET Committee on Computer Research and Applications provides an appropriate forum for coordination of Federal agency programs. Specifically:

- The subcommittee on Computer Networking, Infrastructure, and Digital Communications will develop a coordinated implementation plan for the national research network.
- The subcommittee on Science and Engineering Computing will review the *grand challenges* through the use of high performance computing systems, including the research that will be involved.
- The subcommittee on Computer Research and Development will review the need for advanced software, algorithms, and hardware for future high performance computing systems.

All of the subcommittees will consider appropriate action to secure a foundation of basic research and human resources. In all three subcommittees we expect some overlap of responsibility and interchange of ideas to be compatible with success.

As has been firmly stated, the full cooperation through a shared vision between government, industry and the research community will be a necessary ingredient for the successful implementation of this strategy. The FCCSET Committee on Computer Research and Applications therefore calls for timely consideration of the vision and strategy by representative bodies of the research community and industry.

It is essential, however, that implementation of the strategy be undertaken in a timely manner. There is a need to follow through on the breakthroughs that occurred partially as a result of federal investment in the early 1980s. The fast pace of development dictates that appropriate Federal efforts are needed to help ensure continued excellence in high speed networking technology and leadership in high performance computing. Foreign investment in technology development in these key areas has increased dramatically. The prudent strategy is to maintain a consistent strong lead in research and to transfer the results as quickly as possible to American industry.



## COST ESTIMATES

Many of the basic elements of the high performance computing strategy are already being implemented as part of ongoing agency programs at DOE, DARPA, NSF, NASA, and other Federal agencies, and important progress is being made. The FCCSET Committee activity has contributed to achieving a shared vision, and early coordination is already occurring in anticipation of implementation of the strategy. Implementation of the strategy involves three principal funding components, including the national research network, joint research to address the "grand challenges," and basic research in high performance computing architecture, custom hardware design, software, algorithms, and supporting technologies. Multiple agencies are involved in the implementation and funding of each of these components.

The funds that would be associated with each of these components are described below. Obviously, any incremental funding must be evaluated and approved within the context of current activities and research needs in other high priority fields. Currently, the Federal government is spending about \$500M per year on all aspects of high

**Summary of Additional Funds**  
(Millions of Dollars)

Current Funds		Yr 1	Yr 2	Yr 3	Yr 4	Yr 5
50 <sup>a</sup>	National Research Network	5	5	5	0	0
	Stage 1	5	5	55	55	55
	Stage 2	40	40	40	40	40
	Stage 3					
150 <sup>b</sup>	Joint Research in Computational Science and Engineering	30	60	90	120	150
300	Basic Research in Computer Science and High Performance Computing	60	120	180	240	300
500	<b>TOTAL (above current funds)</b>	140	230	370	455	545
	<b>Funding Increase by Year (noncumulative)</b>	140	90	140	85	90

a Estimated network research and support in grants and contracts.

b Estimated operating costs for existing computational science facilities.

performance computing. Funding for the activities recommended in this report would increase this base by \$140M in additional resources for the first year, growing to an additional \$545M per year in 5 years.

**National Research Network.** Current operating costs for the present collection of research–support networks operated by DARPA, NSF, DOE, and NASA is approximately \$50M per year; the figure is uncertain because many subnetworks are funded by increments on research grants and contracts, rather than being centrally supported. Currently the interconnection of existing agencies’ networks is planned within existing budgets. A significant increase in investment is needed to achieve the required capability. This investment could occur in three concurrent stages.

The *first stage* activity would involve an immediate upgrade to 1.5 Mbit/sec of the existing research–support networks. This would cost \$15M over three years.

The *second stage* would expand upgraded network services (45Mbit/sec) to 200 to 300 research installations, using primarily fiber–optic trunk facilities. Development costs for this stage would be \$5M per year of additional funding. Operation of the upgraded network would commence in three to five years, with operating costs of approximately \$50M per year. Since the transition from the first stage to the second stage network could not be instantaneous, initially the full operating cost of the second stage network would necessitate additional funding; that requirement will diminish to the extent that the first stage network is phased out.

The goal of the *third stage* would be to deliver one to three Gbit/sec to selected research facilities, and 45 Mbit/sec to approximately 1000 research sites. Research and development costs for this project are estimated at \$400M of new funds, spent over ten years; after ten years, operating costs would be about \$200M per year unless some tariff relief is achieved.

**Joint Research in Computational Science and Engineering.** Current operating costs for existing computational science laboratory facilities is approximately \$150M per year. Additional investment would be required to upgrade the existing facilities and/or to establish additional joint research activities, with government, industry, and university participation, to address approximately specific problem areas, including selected *grand challenges*. Many of these joint research efforts will involve multiple physical sites connected by the research network. The investment in these research activities supports pursuit of the grand challenges. This includes personnel to develop computational approaches in terms of theory, algorithms, and software, and the acquisition of modern computing equipment. Estimated Federal costs average \$15M per year to establish and sustain each grand challenge. The joint research activities would be introduced at the rate of two per year. Overall investment will be approximately \$30M per year initially, increasing to \$150M per year in five years as new grand challenges are added.

**Basic Research in Computer Science and High Performance Computing.** Current Federal investment in advanced computer research is estimated at \$300M in FY88. Over the past four years, investment in these areas has grown at 15% per year. The rate of increase appears to be declining, however, at a time when increased investment appears to be needed. Sufficient resources should continue to be allocated to take full advantage of the high performance computing opportunities that now exist including design and prototype development of systems capable of trillions of operations per second. A second important element is stable funding, which is required to preserve the long-term strength of the research community.

Other countries are also devoting considerable resources in this area. For example, the Japanese government supports two projects which directly address supercomputer development: The Fifth Generation Project and the Superspeed Project. Support for each of these is estimated to be in excess of \$100M per year. In addition to this government support, Japanese industry is investing considerably more to develop high performance computers. Japanese government and industry are also investing amounts comparable to those recommended here to develop high bandwidth research networks.

## ACKNOWLEDGMENTS

Office of Science and Technology Policy guidance was provided by Michael Marks. Stephen L. Squires, Defense Advanced Research Projects Agency, acted as Executive Secretary for this report. Technical assistance was provided by William L. Scherlis, Defense Advanced Research Projects Agency; along with Kathleen Bernard, Office of Science and Technology Policy; Charles N. Brownstein, National Science Foundation; Leslie Chow, National Aeronautics and Space Administration; and Michael Crisp, Department of Energy.

# FCCSET COMMITTEE ON COMPUTER RESEARCH AND APPLICATIONS

Paul G. Huray (Chair)  
Office of Science and Technology Policy

## SUBCOMMITTEES

### Science and Engineering Computing

James F. Decker (Chair)  
Department of Energy

James Burrows  
National Bureau of Standards  
John S. Cavallini  
Health and Human Services  
Melvyn Ciment  
National Science Foundation  
John Connolly  
National Science Foundation  
Craig Fields  
Defense Advanced Research  
Projects Agency  
Harlow Freitag  
Supercomputer Research Center  
Randolph Graves  
National Aeronautics and Space  
Administration  
Norman H. Kreisman  
Department of Energy  
Lewis Lipkin  
National Institutes of Health  
Allan T. Mense  
Strategic Defense Initiative Office  
David B. Nelson  
Department of Energy  
C. E. Oliver  
Air Force Weapons Lab  
John P. Riganati  
Supercomputer Research Center  
Paul B. Schneck  
Supercomputer Research Center  
K. Speierman  
National Security Agency

### Computer Research and Development

Saul Amarel (Chair)  
Defense Advanced Research  
Projects Agency

Donald Austin  
Department of Energy  
C. Gordon Bell  
National Science Foundation  
James Burrows  
National Bureau of Standards  
Bernard Chern  
National Science Foundation  
Peter Freeman  
National Science Foundation  
Lee Holcomb  
National Aeronautics and Space  
Administration  
Charles Holland  
Office of Naval Research  
Robert E. Kahn  
Computer Science Technology  
Board  
Daniel R. Masys  
National Institutes of Health  
Robert Polvado  
Central Intelligence Agency  
David Sadoff  
Department of State  
William L. Scherlis  
Defense Advanced Research  
Projects Agency  
K. Speierman  
National Security Agency  
Stephen L. Squires  
Defense Advanced Research  
Projects Agency  
Charles F. Stebbins  
Air Force Systems Command  
Daniel F. Weiner, II  
Joint Tactical Fusion Program

### Computer Networking, Infrastructure and Digital Communications

C. Gordon Bell (Chair)  
National Science Foundation

Ronald Bailey  
National Aeronautics and Space  
Administration  
Sandra Bates  
National Aeronautics and Space  
Administration  
James Burrows  
National Bureau of Standards  
John S. Cavallini  
Health and Human Services  
Thomas Kitchens  
Department of Energy  
James Oberthaler  
National Institutes of Health  
Dennis G. Perry  
Defense Advanced Research  
Projects Agency  
Arnold Pratt  
National Institutes of Health  
Shirley Radack  
National Bureau of Standards  
Rudi F. Saenger  
Naval Research Laboratory  
Daniel VanBelleghem  
National Science Foundation  
Stephen Wolff  
National Science Foundation


- 3 Newslog
- 10 Forum
- 14 Book reviews
- 16 Software review
- 17 The engineer at large
- 18 Innovations
- 19 Program notes
- 22 Video
- 24 Speakout
- 26 Faults & failures
- 58 EEs' tools & toys
- 60 Papers are invited
- 61 Calendar
- 63 New publications
- 67 IEEE tables of contents
- 82 Scanning THE INSTITUTE
- 82 Coming in *Spectrum*

*Cover: Before destroying all intermediate-range missiles under the recently signed arms pact, Soviet and United States workers may salvage the nuclear warheads. This uranium-235 button, the end product of a recovery process at the Oak Ridge Y-12 Plant in Tennessee, may be recast and machined for warhead components. The 9.5-lb button, actually lead-colored, is valued at about \$200 000. Photo: Martin Marietta Energy Systems Inc.*



Field Editor Perry pays for a tank of gas with her bank credit card

46



Minicomputer pioneer Bell lectures on the merits of a national network

54

### STAFF

**Editor and Publisher:** Donald Christiansen  
**Executive Editor:** Edward A. Torrero  
**Administrative Editor:** Ronald K. Jurgen  
**Senior Technical Editor:** Gadi Kaplan  
**Managing Editor:** Mark A. Fischetti  
**Senior Associate Editor:** Trudy E. Bell  
**Associate Editors:** John A. Adam, Elizabeth Corcoran, Karen Fitzgerald, Erin E. Murphy, Gary Stix, John Voelcker, Katherine Wollard, Glenn Zorpette  
**Field Editor (West Coast):** Tekla S. Perry  
**Chief Copy Editor:** Margaret E. Eastman  
**Copy Editor:** Michael Olver  
**Editorial Researcher:** Robert Davidson  
**Contributing Editors:** Evan Herbert, Alexander A. McKenzie, Michael F. Wolff  
**Manuscript Review:** Rita Holland (Supervisor)  
**Design Consultant:** Gus Sauter  
**Art and Production:** Janet Mannheimer (Director)  
**Editorial Production:** Ruth M. Edmiston (Manager), Marcia Meyers (Production Supervisor), Peter Ruffett (Typographer)  
**Art:** Morris Khan (Technical Graphic Artist), Elaine Wesson Da Costa (Manager)  
**Advertising Director and Associate Publisher:** William R. Saunders  
**Administrative Assistant:** Wendy Tinkler  
**Research:** Hendrik Prins (Manager)  
**Circulation and Research Administrator:** Joanne Biley  
**Mail List Sales:** Geoffrey L. Klapisch (Manager)  
**Business Manager:** Carl Maier  
**Marketing:** Robert D. Moran (Manager)  
**Administrative Assistant to the Editor and Publisher:** Nancy T. Hantman

### ADVISORY BOARD

**Chairman:** Jerome J. Suran  
 Edward C. Bertnoff, Richard W. Damon, N. Rex Dixon, Thelma Estrin, Ronald G. Hoelzeman, Robert W. Lucky, G.P. Rodrigue, Michael J. Whitelaw


### EDITORIAL BOARD

**Chairman:** Donald Christiansen  
 Wallace B. Behnke, Peter D. Blair, James E. Brittain, S.J. Campanella, Walter S. Ciciora, Theo A.C.M. Claasen, Almon H. Clegg, Karsten E. Drangeid, Lester Eastman, Kimberly D. Emerson, Ernst Feldtkeller, Bob L. Gregory, Frederick Hayes-Roth, Charles K. Kao, Fred L. Katzmann, John L. McLucas, M. Granger Morgan, Eric Nussbaum, Alfred R. Potvin, Howard Sachs, John B. Slaughter, Cary Spitzer, Michiyuki Uenohara, Thomas P. Wolf

Postal identification statement is on page 6.

# Gordon Bell calls for a U.S. research network

*Advances in computing and communications research in the United States will depend on a powerful data-communications network, perhaps best provided by the Federal government*



If a research cardiologist at Boston University Medical Center urgently needs to review cardiac images with a colleague at the Mayo Clinic in Minnesota, he can either express-mail his material to the clinic or fly there. What he cannot do is transmit the material instantaneously from the computer workstation in his office. On the other hand, researchers at the Massachusetts Institute of Technology do not

lack networks. They can communicate with many research organizations around the world, but they must use the right one of a dozen networks to do it.

These scenarios point up just two absurdities of the present situation in U.S. computer networking. Existing networks not only lag well behind the growing needs of the research community—they are too fragmented to develop unaided into a single, coherent system.

The most viable solution is a national research network organized and maintained by the Federal government. Access to information has never been more important than it is today, and the ability to fully exploit information resources—be they individual researchers, research teams, databases or supercomputers—will determine how competitive any group or nation is. Any new proposal costs, of course. But a single national network, jointly supported by all the Government agencies now running independent networks, could well save money.


## *Quantity does not equal quality*

Computer networking among scattered facilities in the United States began in 1969, when the Defense Advanced Research Projects Agency (Darpa) established Arpanet. The network started out as a means of sharing expensive equipment, databases, files, and above all time on computers that would otherwise have lain idle. What developed, however, was a completely different style of interaction. Utilizing the ability to send mail and large documents electronically, researchers have built electronic bulletin boards and held extensive forums and conferences.

Today's Arpanet is conceptually identical to the network of the early 1970s. But it can do little more than swap computer-mail messages now that the number of machines has mushroomed beyond a hundred switching computers that connect hundreds more shared computers and workstations. The network could be upgraded, with great difficulty. But the Defense Communications Agency, which oversees it, is reluctant to run a civilian research network.

Since Arpanet was established, some 36 other major research networks have sprung up around the world, all based on variations of Arpanet's method for packet-switching data. Typically

*C. Gordon Bell Ardent Computer Corp.*



*If science is willing to wait, a national network just might eventually evolve, according to Bell. But patience seldom assists progress.*

Some are addicted to information as predigested experience, a cognitive fast food. This is a danger for the new generation, raised on the instant knowledge of television and the computer. They need discipline to filter and edit useful information, and beyond this, to develop deeper interest and understanding of the world and of themselves.

Some self-developers feel their possibilities are unlimited and so are unrealistic about what they can achieve. They seem like neglected, hyperactive children as they flit compulsively from one "developmental experience" to another.

Commitment and intimacy are problems for self-developers. "How much of myself am I willing to commit?" is a popular phrase for this new generation. Of the self-developers we interviewed, a high percentage (25 percent) have been divorced and fewer than half are in their first marriage (48 percent). In contrast, only 6 percent ever divorced in *The Gamesman* sample of high-tech managers. Given that self-developers are on the average younger than the other types, the difference may become even greater in the future, as more self-developer marriages break up.

### *Motivating the self-developer*

To motivate the new generation, managers must give them opportunities to develop marketable skills. Self-developers are also stirred by the chance to expand their knowledge, improve their well-being, live a more enriching life through travel abroad.

They see money as only part of a job's reward and, as long as their basic financial needs are met, weigh income against other payoffs like the opportunity to develop skills, time off, health care, child care, exercise facilities, and a friendly atmosphere.

When managers fail to make work meaningful to self-developers, the self-developers will find ways to express themselves outside the company—or will quit. In fact, they switch jobs more easily than traditional professionals. This willingness to quit an oppressive job, the dual-career family, money in the bank, and secondary entrepreneurial ventures, all help to fan the spirit of independence.

Self-developers find it relatively easy to jump ship because, they say, work isn't the only important thing in their lives. The ones with families want to keep a balance. A highly competent and motivated executive, aged 38, says: "I work 50 hours a week. I come in early and go home late, but I leave it at the office. I am not going to push my little son off my lap because he is messing up papers from work."

For self-developers, physical well-being also requires attention. A 29-year-old manager of information systems says: "We watch the things we eat. We run. If you're a workaholic, after a while you're not really productive."

### *The move into management*

While self-developers are critical of the managerial hierarchy, they will move up, albeit reluctantly: it is safer and crisper to be a professional self-developer than to commit oneself to risky projects and the education of others. But they recognize they need power to get things done.

The good news is that self-developers are natural facilitators. They can create an open atmosphere where views are exchanged, conflict becomes constructive dialogue and study, and consensus is achieved. They can facilitate well because they are egalitarian and interested in other's views and ideas.

All self-developer managers believe that to succeed they must create a motivated team. For them, being a team player does not mean group thinking. It means playing a special role on a team where each player has a say in how to implement strategy.

Self-developers succeed best as managers when they institute good practices: frequent evaluations, team meetings, and training in group process and problem-solving.

The weaknesses of self-developers as managers stems from their reluctance to make commitments, which appears to me a significant problem in about 40 to 50 percent. And no one gains authority in the minds of others without commitment to projects and

### **The tellers of the tale**

The material in this article is excerpted from *Why Work*, to be published in March by Simon & Schuster, New York. While my earlier book, *The Gamesman*, focused on the elite, high-technology company, the study upon which *Why Work* is based was much broader. Its conclusions apply to engineers as well as to other professionals in today's technoservice economy.

*Why Work* is the result of seven years of research involving nine companies and more than a dozen Federal, state, and municipal agencies. Over 350 people at all levels, from chief executives to front-line service employees, were interviewed about their work values in sessions lasting from 1½ to 3 hours. They were asked what satisfied and dissatisfied them at work, how they defined service, how they wanted to be managed, how they managed others, how they relate to customers, clients, and co-workers, and about their family background and goals at work. Many of the questions were similar to those used by *Spectrum* in its series of articles by Tekla S. Perry from December 1983 to May 1984 on engineering environments.

Businesses where employees were interviewed and surveyed include AT&T, U.S. West, a large big-eight accounting firm, an innovative insurance company, a TV broadcasting company, a large supermarket chain, the service division of a large oil company, a company producing information systems, and Scandinavian Airline Systems. The government agencies include the Internal Revenue Service; the Commerce Department; Action (which comprises the Peace Corps, Vista, and the Older Americans Volunteer Program); the Federal Aviation Administration; the National Aeronautics and Space Administration; the departments of agriculture, justice, and defense; the Federal Trade Commission; the Veterans Administration; the National Highway Traffic Commission; the Library of Congress; a statewide health department; two hospitals; a county health department; a city tax office; a social worker office in California; and a municipal library. Besides these, the study drew on interviews conducted by my colleagues and students with police officers in two metropolitan departments, entrepreneurs in the U.S. and Sweden, and middle managers in Japanese banks and trading companies.

—M.M.

people. If self-developers are going to become effective leaders, they must find meaning in caring for others and taking responsibility for larger enterprises.

### *To probe further*

The motivation of entrepreneurs and managers in high-technology industry is described in Michael Maccoby's best-seller, *The Gamesman* (Simon & Schuster, New York, 1977). This book originated as a *Spectrum* article, "Winning and losing at work" [July 1973, pp. 39-48].

A new model of leadership for the 1980s that combined improved competitiveness and the quality of working life is described in Maccoby's book, *The Leader* (Simon & Schuster, 1981).

The relationship of employees of all levels in major corporations and government agencies to their workplace is analyzed in *Why Work*, a book by Maccoby to be published in March by Simon & Schuster.

### *About the author*

Michael Maccoby directs the Project on Technology, Work and Character, a center for research and consulting in Washington, D.C., and he also acts as consultant to business, government, and labor unions. He has a Ph.D. in social relations (1960) from Harvard University. ♦

these networks convey data at rates of 1.2 to 56 kilobits per second, using incompatible communication protocols.

Federal agencies usually support an average of two independent wide-area networks. Often these networks go to different buildings on the same site—be it university campus or Federal laboratory—wasting resources. Yet the Government cannot even begin to estimate the current costs because each Federal agency considers its networking expenses proprietary information.

In any event, more networks do not automatically translate into greater capabilities. The situation is reminiscent of telephone systems in the early 1900s, when a town might support several distinct company telephone networks, forcing subscribers to use a deskful of phones. Theodore N. Vail, president of the American Telephone & Telegraph Co., however, successfully corralled the local companies under the banner: "One policy, one system, universal service." Similarly today, easily a dozen incompatible networks may overrun just one university campus.

Their incompatibility, if not their numbers, is fortunately waning. Within the past two years, most networks have begun migrating toward Darpa's Transmission Control Protocol/Internet Protocol (TCP/IP) standards and are committed to using the internationally approved Open Systems Interconnect (OSI) standards as they become available. Identical protocols for exchanging data not only make it easy to connect networks but to have them share common links to save equipment costs.

On university campuses, researchers are wiring their array of local-area networks (LANs) into campus area networks (CANs). But at present, most campus area networks are isolated. Since LANs typically operate at 10 megabits per second and CANs operate at up to 100 Mbits/s, very fast wide-area networks will be required to connect these CANs into a global network.

### *Faster computers, fiercer demands*

The rise in computer speed on campuses has triggered other developments. It has revolutionized the nature of the data that researchers want to share, boosted the demands on supercomputers, and encouraged new forms of collaborative research.

A report released last July from the National Science Foundation (NSF) stressed how sorely researchers need networking to create visual and hence more comprehensible representations of supercomputer output and to exchange high-quality graphical data, including photographs. Since a high-resolution workstation displays about a million pixels that each can change as often as 10 to 60 times a second, the bandwidth required for sending dynamic pictures varies from 10 to 60 Mbits/s for a black and white display. For a color display, it soars to 320 to 1920 Mbits/s.

These predictions assume data is not compressed before being transmitted. Depending on the application, data compression techniques can reduce the necessary bandwidth by a factor of 10 to 1000. Nevertheless, the data requirements of a national research network still exceed current capabilities.

Connecting several hundred workstations and high-performance computers would require a network capable of delivering hundreds of megabits per second. Visualizing mechanical parts, medical images, and geological data can demand the transfer of some 4 gigabytes of data among workstations or from supercomputers to workstations. Without data encoding, a 45-Mbit/s link would transmit those 4 Gbytes within 10 minutes. Today's networks are over 1000 times slower!

Remote access to supercomputers is particularly critical since few institutions can afford their own. As supercomputing time is limited, researchers usually run pieces of a program on smaller computers or workstations, then transmit the entire program and database over a network to a supercomputer. Bandwidth requirements for such transmissions easily exceed 1 Mbit/s.

Increasingly, the problems under study require the active collaboration of researchers who are scattered among various research institutions. Collaboration technology—an emerging area in which researchers work together over a network—depends on a range of networking capabilities, including: compound docu-

ment transfer and the simultaneous viewing and editing of documents that combine text, graphics, pictures, and voice; computer conferencing with attendees interacting through both pictures and conversation; design reviews performed on a common document; and the ability to remotely control and interact with special laboratory and industrial facilities.

Clearly, huge volumes of data will pour through the national research network. Other countries are already acting: the Japanese and European governments are busy building fiber-optic computer networks that will transmit gigabits of data per second. In Japan, most major research centers already plug into high-speed networks that enable them to store and transmit international scientific and engineering material.

### *Spontaneous generation unlikely*

All these developments point to the need for a national research network in the United States. But might that network evolve over time, without a centrally administered plan?

Some argue that a sort of national network may emerge with advances in fiber-optic technology. Indeed, today's fiber-optic communication links promise 1000 to 100 000 times the capacity and speed of traditional cable and satellite channels. But the cost of creating a large enough fiber-optic network remains prohibitive.

Supply and demand play no part in the price of fiber-optic networks. A vast amount of fiber remains unused. But the prices of using fiber-optic links are primarily based on the rates for transmitting voice communication on coaxial cables. Since a fiber-optic bundle can transmit several orders of magnitude more data than a coaxial cable, the price per bit transmitted should be quite low. Carriers may blame regulatory agencies for the voice-based pricing, but the situation almost forces any organization wishing to send data in this way to set up a private network.

Furthermore, switching equipment to exploit even the DS3 standard of 45 Mbits/s is unavailable. Researchers have developed some designs and prototypes for 45-Mbit/s packet switches. But no companies currently manufacture such products. Suppliers of communications services are loath to invest in fast data-communications networks, since they already have large voice-traffic networks and many customers. Proposals for national standards for higher speeds are moving at glacial rates.

The integrated-services digital network (ISDN) is often touted as a panacea. But it has moved much too slowly to hold much promise for a national network in the 1990s. Local and international carriers are still thrashing out technical specifications for compatibility. The regional Bell operating companies are not especially cooperating with each other to set up ISDN standards. There are barely standards for low-level protocols, and there are no standards at higher levels. Moreover, when high-speed fibers do terminate in switching offices, distributing them to local users takes an inordinate amount of time and effort—the so-called "last mile" problem. In fact, U.S. manufacturers are losing the ability, if they ever had it, to build ISDN equipment since they buy the hardware from their international partners. Wisconsin Bell collaborates with Siemens. Pacific Northwest Bell works with Northern Telecom. Mountain Bell, AT&T, GTE, Illinois Bell, NYNEX, and Southern Bell all have various links to Ericsson, NEC, Northern Telecom, Siemens, and Fujitsu.

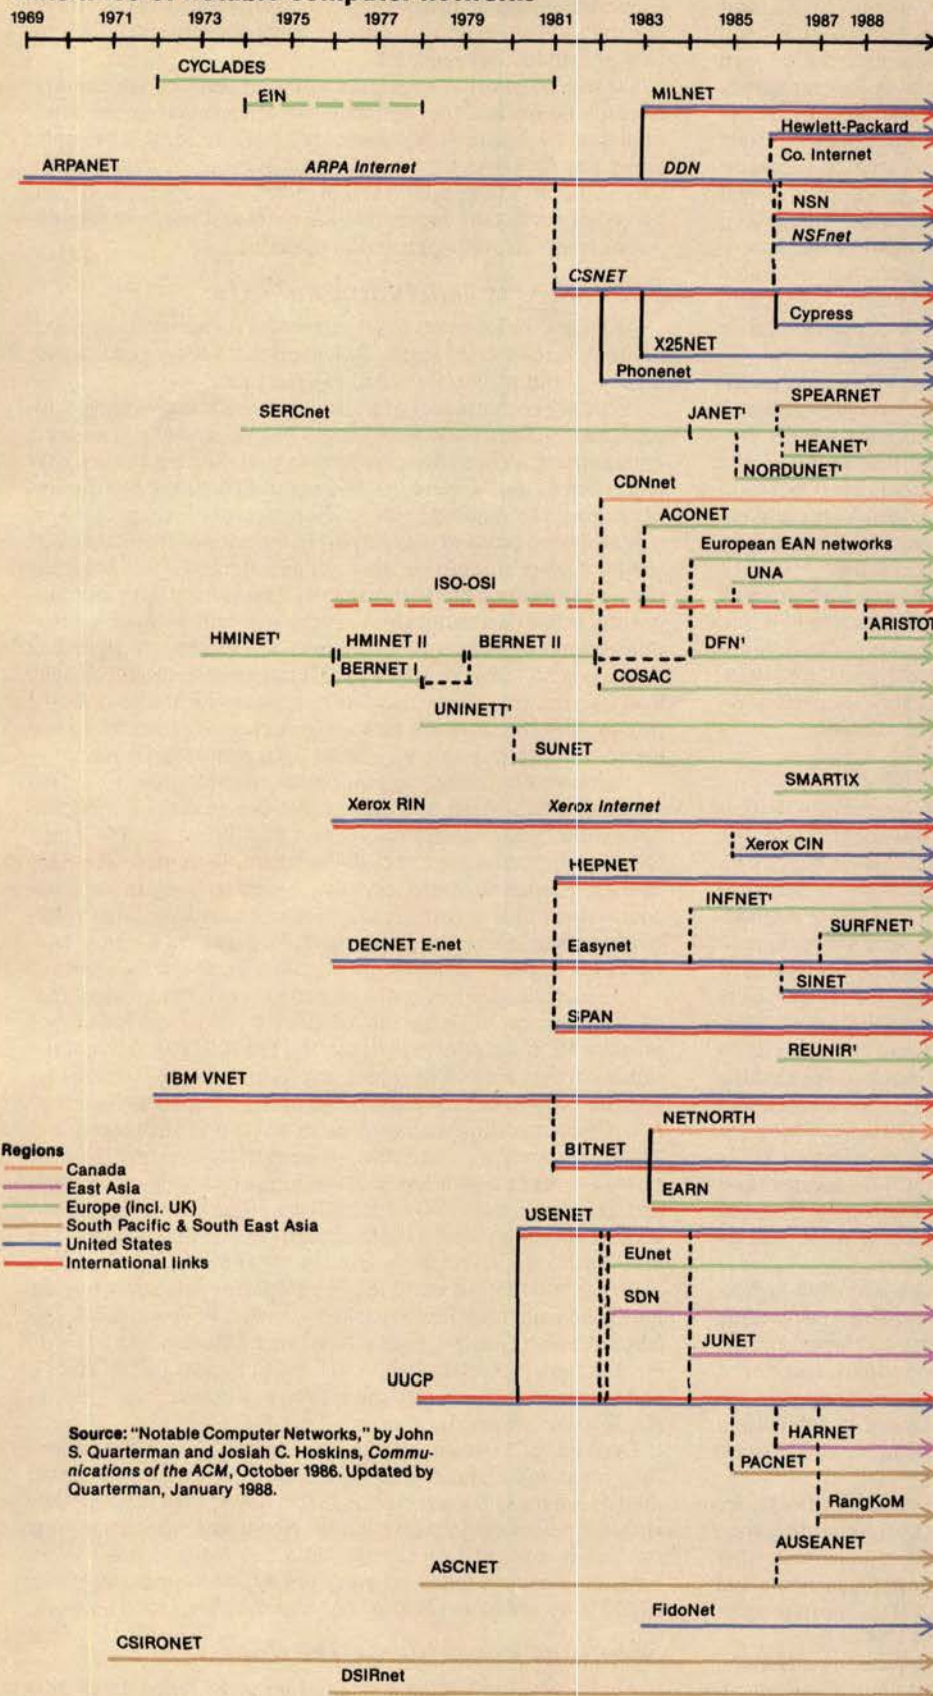
Defined ISDN protocols remain a factor of 200 to 2000 slower than local and campus area networks, a factor of 1000 slower than 45 Mbits/s, and a factor of 30 000 slower than is potentially viable for fiber-optic transmission. Nor is any work under way for higher transmission speeds. Yet a link for 45-Mbit/s transmission is needed today on every university campus. In effect, ISDN is irrelevant to the needs of a national research network.

### *Solutions knocking on the door*

The closest the United States has come so far to developing a national research network is the 56-kbit/s NSFNET. Last November, the NSF began to extend the NSFNET backbone beyond the present six supercomputing centers to include seven



# Time lines of notable computer networks



- Regions**
- Canada
  - East Asia
  - Europe (incl. UK)
  - South Pacific & South East Asia
  - United States
  - International links

Source: "Notable Computer Networks," by John S. Quarterman and Josiah C. Hoskins, *Communications of the ACM*, October 1986. Updated by Quarterman, January 1988.

- ACONET: Akademisches Computer Netz
- ARISTOTE: Association de Réseaux Informatiques en Système Totalement Ouvert et Très Elaboré
- ARPANET: Advanced Research Projects Agency Network
- ASCNET: Australian Computer Science Network
- AUSEANET: Austroasian Network
- BERNET I, II: Berlin Network
- BITNET: Because It's Time Network
- COSAC: Communications sans Connexions
- CSIRONET: Commonwealth Scientific & Industrial Research Organization Network
- CSNET: Computer Science Network
- DDN: Defense Data Network
- DECNET E-net: DEC Engineering Network
- DFN: Deutsche Forschungsnetz
- DSIRnet: Government network in New Zealand
- EARN: European Academic Research Network
- EIN: European Informatics Network
- EUnet: European Unix Network
- HARNET: Hong Kong Academic & Research Network
- HEANET: Higher Education Authority Network
- HEPNET: High Energy Physics Network
- HMINET I, II: Hahn-Meitner Institut
- IBM VNET: Virtual Network
- INFNET: Istituto Nazionale Fisica Nucleare
- ISO-OSI: International Organization for Standardization-Open Systems Interconnect
- JANET: Joint Academic Network
- JUNET: Japanese Unix Network
- MILNET: Military Network
- NORDUNET: Nordic University Network
- NSFnet: National Science Foundation Network
- NSN: NASA Science Network
- PACNET: Pacific Network
- RangKoM: Rangkaian Komputer Malaysia
- REUNIR: Réseaux des Universités et de la Recherche
- SDN: System Development Network
- SERCnet: Science Engineering Research Council Network
- SINET: Schlumberger Information Network
- SPAN: Space Physics Analysis Network
- SPEARNET: South Pacific Education and Research Network
- SUNET: Swedish University Network
- SURFNET: Dutch university network
- UNA: Universitäts-Netz Austria
- UNINETT: Nordic university network
- USENET: Users' Network
- UUCP: Unix to Unix Copy
- Xerox CIN: Corporate Internet
- Xerox RIN: Research Internet

Note: CDNet, CYCLADES, EAN, NETNORTH, and SMARTIX are not acronyms.

The past five years have seen the number of networks soar dramatically. Many initially used transmission protocols and technology developed by one or more older networks; these are indicated by the vertical dashed lines connecting networks. (More recently, some networks have begun to use other protocols, particularly ISO-OSI and Arpanet standards.) Solid vertical lines between networks indicate systems under closely related administrations. Dashed horizontal lines indicate protocols or demonstration systems, rather than operational networks. Networks in italics are internets—several networks tied together that use the same transmission protocols.

<sup>1</sup>Participants in the Réseaux Associés pour la Recherche Européenne network (RARE)

regional, university-based research networks. The backbone should be running at 1.5 Mb/s by July under the management of Merit Inc., which is based at the University of Michigan, Ann Arbor, and assisted by IBM Corp. and MCI Communications Corp. Unfortunately, the fiscal 1988 budget allocated by Congress for NSFNET barely keeps the network alive.

But the NSF is only one agency. It has no authority for incorporating other Federal networks, no timetable for upgrading to 45-Mbit/s rates, and certainly no budget for doing so.

The U.S. government is slowly recognizing the need for a national research network. In 1986, Congress requested that the Office of Science and Technology Policy (OSTP) study the problems and options of developing a communications network for research computers, including supercomputers at U.S. universities and Federal research facilities, and provide a plan for action by August 1987. The OSTP accordingly established a new inter-agency group, the Federal Coordinating Council for Science, Engineering, and Technology for Computer Research and Applications. The council finished a three-volume report in on time; the OSTP finally sent a summary of the report to Congress late last November.

I chaired the subcommittee on computer networking, infrastructure, and digital communications. In the report, we strongly urged that the Government create a national research network to "foster and enhance the U.S. position of world leadership in computer networking." I believe the situation is far worse; we have already lost leadership in this field. By developing a network that enables U.S. researchers at all universities, national labs, and companies to share resources and ideas, the country just might regain its footing.

Implementing the network can be done in three steps. Stage 1 should be to connect Arpanet with other networks supported by Federal agencies over the next two years. If coordinated and centrally managed, these facilities could unite many computer networks into a seemingly single computer network. Operating the backbone and major regional networks at 1.5 Mb/s should open up a whole new set of library and educational services.

The Government should provide funds for stage 1. The annual cost for such an upgraded service is likely to be \$5 million and should be shared by the five Federal agencies that support the most networking: NSF, Darpa, the Department of Energy, the National Aeronautics and Space Administration, and the Department of Health and Human Services. As part of stage 1, a network manager should be selected and made responsible for upgrading the network speed from 1.5 to 45 Mb/s within three years (stage 2) and to many gigabits a second by the late 1990s (stage 3).

Stage 2 should include upgrading and expanding the existing facilities at 200 to 400 U.S. research institutions at data communication rates of 1.5 Mb/s, or T1 rates. This work would require new funding at approximately \$5 million per year over five years. The estimate assumes that the price of T1 lines will halve over the next five years—a modest assumption, since oversupply pushed prices down more than that in the second half of 1987 alone. Operating expenses for the upgraded facilities are likely to be \$50 million annually. While Government should support the first years, eventually users should cover the costs of the network service, the same way they now pay for telephone service.

Establishing a vigorous, focused program of network research and development is critical to stage 3. Some \$400 million would be needed over 10 years to advance networking technology and make it possible to transmit and switch 3 Gb/s by the early 2000s. Such a network would have 100 000 times more capacity than those currently available and enable researchers to communicate instantaneously.

### Who will take the lead?

The Government is certainly not the only hope in this situation. Any one or a combination of the existing telecommunications suppliers could pre-empt Federal efforts to build a nation-

al research network by simply building the network and offering the service for sale. A highly aggressive, imaginative telecommunications industry could view the network as the major, large-scale social experiment of this century and far-sighted preparation for the next.

Achieving stage 2 (45 Mb/s) in three years would take only a small fraction of this industry's research, development, and operations budget. Government should encourage such efforts in every way possible, but, to judge from history, the country cannot depend on such an initiative.

The OSTP report took a different tack. It recommended appointing a lead agency to oversee networking. In January, NSF volunteered for the job; the agencies that participated in the OSTP report agreed. Still, there are no examples of a single agency supplying a facility for the entire research enterprise. In fact, agency behavior is byzantine: each wants its own facility, no matter what the cost. Both the executive and legislative branches of the Federal government are simply incapable of setting up a line item to be funded and administered either by an agency or by an inter-agency group, even though the facility would support the entire research and higher education community.

A radical approach, though, could work: select a private-sector company to manage and develop the network, and provide it with a budget, to which every agency would contribute under NSF guidance. (Each would list its support for the network as a single line item in its budget. Each would also relinquish control of its networks to the manager.) Assure the network manager of steady support—both fiscal and political—for the first five operating years or so of the network. And instruct the manager to devise a plan for gradually shifting all operating costs to users.

Both the national research network and supercomputer facilities could be funded in this fashion. For lack of a common facilities budget, the Federal agencies at present have no choice but to fund and build their own, inevitably overlapping networks and supercomputer centers. Perpetuating this situation only wastes more dollars and more time.

Building this network is not a difficult problem for the U.S. engineering community. But the United States lacks leadership in communications as well as anything resembling a coordinated Federal science and technology policy. Our best hope may be the research community, if it can successfully mobilize its resources. After all, it stands to benefit the most and the soonest from a national research network, even if the country's strength and prosperity over the long haul is what is ultimately at stake.

### To probe further

The summary report, "A Research and Development Strategy for High-Performance Computing," sponsored by the Office of Science and Technology Policy, includes an outline for a national research network. Copies are available from the OSTP, attn. Kathleen Bernard, New Executive Office Building, Room 5005, Washington, D.C. 20506. Thoroughgoing descriptions of many existing research networks are given in "Notable Computer Networks," by John S. Quarterman and Josiah C. Hoskins, *Communications of the ACM*, vol. 29, no. 10, Oct. 1986, pp. 932-968. *IEEE Communications Magazine* has also carried many articles on computer networking. Among them, "Research computer networks and their interconnection," by L.H. Landweber, D.M. Jennings, and I. Fuchs, June 1986, pp. 5-17, is a good introduction.

### About the author

Best known for his work on the VAX superminicomputer, C. Gordon Bell [F] is vice president for research and development at Ardent Computer Corp., Sunnyvale, Calif. (formerly Dana Computer). Under the Federal Coordinating Council for Science, Engineering, and Technology, Bell chaired the subcommittee on computer networking, infrastructure and digital communications. In 1986 and 1987, he served as assistant director for computing at the National Science Foundation. He earned his B.S. and M.S. degrees in electrical engineering at MIT. ♦

# The National Research Network: The Superhighway of the Computer Age

## Introduction

### Technology Background and the Need for A National Research Network

The benefits of computers networking for scientific and engineering research have been recognized almost from the time that computers became essential research tools in the late 1960's. With the innovation of timesharing, giving multiple users simultaneous access to single machines, then with the use of the telephone network for access from remote locations, came recognition of the economies of resource sharing and the productivity-enhancing effects of improved communications.

ARPAnet, introduced in the early 1970's based on sending packets of information from computer to computer was invented to interconnect the computers using a separate network of special, small packet switching computers. Today, some thirty-six major national and international networks exist (CACM 1986) using either the ARPAnet packet switching ideas, or a modified form of packet switching which is carried out directly by each host computer. State and regional networks are proliferating. NSFNET, an "internet" designed initially to improve access to supercomputer centers, has in the space of two years, forged links among 17 state, regional, and federal agency networks.

The growing base of researchers and educators who rely on networks is beginning to seriously congest the current "system", degrading performance at the very time that demand is growing for ways to link distributed educational and research resources. One result is that enormous effort is devoted to working around problems. NSF's experience with providing access to supercomputers is a case in point.

In the early '80's the lack of access to supercomputing power by the scientific and engineering community caused the formation of the NSF Office of Advanced Scientific Computing, which funded five centers for supercomputers. Given the highly distributed location of users, the need for a national wide area network for computer access and for the interchange of associated scientific information (mail, files, databases, etc.) became clear. It immediately became obvious that existing agency networks both lacked the inherent capacity and were overloaded. In fact, the current federal wide area networks still operate at the speed (56 Kbits/sec) of the original ARPAnet (circa 1972), despite the fact that computers have both increased in number (by a factor of 100) and speed (by a factor of 15). Moreover, the very nature of what is communicated has changed with the power of computers. Research databases are today very large objects. Computer simulations of physical process yield output expressed as three dimensional dynamic process. Computer-based research tools are becoming more interactive, and the problems under study often require the active collaboration of researchers who are distributed in various research institutions. Network technology has fallen seriously behind computing innovations and the subsequent requirements of the research enterprise.

Today's fiber optic communications links offer the ability for a factor of 1000-100,000 increase in capacity and speed over traditional cable and satellite channels. The price of fiber optic links are based on equivalent voice grade circuits, and as such remain high, despite a vast amount of unused fiber. (maybe this should be moved or expanded-it just hangs here)

Each federal agency has an average of two, independent wide area networks that both couple researchers to one another for mail, collaboration, file transfer, etc. and to large, central systems. Each of these networks are likely to go to the same campus (an academic institution, federal or industrial lab), but to a different building. While a few years ago, all wide area networks used different protocols and could not communicate with

one another, today most networks are migrating toward DARPA's TCP/IP protocol with a commitment to use the ISO protocols when available. By having identical protocols, networks can be interconnected together easily (internetting) and share common links to save costs.

Today, campus based local area nets are rapidly becoming standard parts of the local educational and research infrastructure (along with libraries and computing), in all of the nation's institutions of higher education. These Campus Area Networks form networked islands which now require a much higher degree of wide and even global area networking in order to communicate with other institutions. Campuses are being wired to interconnect the array of Local Area Networks- LANs (a wiring scheme for a single building or small cluster of buildings) to form Campus Area Networks - CANS. Since LAN's operate at 10 Mbits/sec and CANS operate at up to 100 Mbits/sec. a very fast wide area network is required in order to interconnect the CANS (and LANs) to form what will become a Global, Local Area Network.

### **Congressional Action and Executive Branch Response**

On June 1986, Senator Gore introduced the Supercomputer Network Study Act of 1986. Public Law 99-383, passed June 21, 1986 by the 99th Congress charged the Office of Science and Technology Policy (OSTP) with conducting "a study of critical problems and current and future options regarding communications networks for research computers, including supercomputers at universities and federal research facilities in the United States".

At OSTP's direction, an interagency group under the auspices of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) for Computer Research and Applications in the Subcommittee on Networking, Infrastructure and Digital Communications formed to carry out the study of the following issues:

- the networking needs of the nation's academic and federal research computer programs, including supercomputer programs, over the next 15 years, including requirements in terms of volume of data, reliability of transmission, software compatibility, graphics capabilities, and transmission security;
- the benefits and opportunities that an improved computer network would offer for electronic mail, file transfer, and remote access and communications; and
- the networking options available for linking academic and research computer, including supercomputers, with a particular emphasis on fiber optics.

The Computer Network Study Planning Group comprised of the Subcommittee members from DoD (including DARPA), DoE, HHS, NASA, NBS, and NSF produced and delivered a three volume report for OSTP in early August 1987 consisting of:

- Volume I. Recommendations to Congress for the National Research Network
- Volume II. A workshop involving 100 researchers, network users, suppliers, and policy officials on the network giving requirements and alternatives, special requirements for supercomputers, internet concepts, future standards and service requirements, computer security, and the government role.
- Volume III. Background papers on internet concepts, network requirements by various agencies, and government's role

### **Summary of Recommendations to OSTP (for Congress) for the National Research Network**

The Subcommittee on Networking, Infrastructure and Digital Communications (NIDC) recommended the following:

"The U.S. should undertake, as a national goal, the establishment of a National Research Network in a staged approach that supports the upgrade of current facilities, and development of needed new capabilities. Achievement of this goal would foster and enhance the U.S. position of world leadership in computer networking.

As rapidly as feasible, the National Research Network should be designed, deployed and maintained as an advanced computer network. This network should interconnect substantially every academic, industrial, and government research establishment and unique scientific resource to encourage scientific collaboration unhindered by distance and to permit the sharing of unique research facilities and resources. Since security of the network is a vital concern, appropriate policies should be adopted to protect the information in the network from threats, vulnerabilities and risks, and to assure a uniform level of security.

Until the National Research Network can replace the current system, existing networks should be maintained and modified as they join the national network. Since supercomputer systems comprise a special and valuable national research resource with very high performance requirements, the responsibility for network access to supercomputers should be vested in the supercomputer centers themselves until the advanced computer network, capable of offering the requisite service level, is operational.

Industry should be encouraged through special incentives to participate in research, development, and deployment of the National Research Network. Tariff schedules which have been set for voice transmission should be re-examined in light of the requirements for transmission of data through computer networking.

To meet the goal for the National Research Network and to set and agenda for the future the following actions are recommended:

- The Subcommittee on Computer Networking, Infrastructure and Digital Communications which was established by the Office of Science and Technology Policy on May 15, 1987, should oversee the first stage in development and operation of the National Research Network -- a coordinated internetwork that would include the Federal agencies that currently operate research supporting networking.
- The FCCSET Subcommittee on Computer Networking, Infrastructure and Digital Communications should identify a lead agency which would be responsible for requesting funds for the National Research Network, and eventually for selecting a contractor to manage the network. The manager would be responsible for Stages 2 (45 Mbit/sec links) and Stage 3 (multi-Gigabit/sec links).
- As a first stage in the development of the National Research Network, the current Internet system developed by DARPA and networks supported by agencies should be interconnected over the next two years. These facilities, if coordinated and centrally managed, have the capability to interconnect many computer networks into a single virtual computer network. The Federal government should encourage and assist research facilities and academic institutions to establish local and campus area networks to connect to the Internet systems. The estimated cost for this proposed upgraded service is \$5 million per year, and should be implemented through the shared resources of NSF, DoE, DARPA, NASA, and HHS.
- In the second stage, new funding for development should be requested at \$5 million per year over the next five years to upgrade and expand the nation's existing computer networks, which support research programs, to achieve data communications at 1.5 Mbits/sec to 200-400 U.S. research institutions. It is estimated that these expanded and upgraded facilities will require an additional funding of approximately \$50,000,000 to operate (GB: assumed the price of T1 lines would decline by a factor of 2 over the next five years, whereas, in the six months since the report, line charges for many T1 lines have already dropped more than a factor of two due to oversupply.)



- In the third stage, a vigorous focused program of research and development for the National Research Network should be immediately established. A total of \$400 million is needed over ten years to advance the knowledge base and technology of computer network capabilities in order to achieve data communications and switching capabilities to support transmission of three billion bits per second within fifteen years. This will support a network 100,000 times more capable than currently available and will be essential to foster scientific collaboration and sharing of research resources. When fully deployed, the cost of operating this advanced network is estimated to be \$400 million per year, given the current commercial tariffs for data communications.

Support should be given to the development of standards and their harmonization in the international arena. Aggressive action is needed to increase user participation in the standards development process, to get requirements for standards expressed early in the development process, and to speed the implementation of standards in commercial off-the-shelf products. It is essential that standards development be carried out within the framework of overall systems requirements to achieve interoperability, common user interfaces to systems and enhanced security."

#### **Motivation for the National Research Network ... The Uses**

Based on the experience of the ARPANet, it is very difficult to predict the ultimate use of National Research Network. The network was initially justified on the basis of being able to share facilities including particular programs, databases, and files including being able to absorb idle computer capacity. What happened was that a completely different style of interaction developed based on being able to send mail and large documents. Finally, extensive public forums and conferences are held through bulletin board and computer conferencing. Remote terminal access is negligible, and file interchange was relatively smaller than expected.

With supercomputing both remote access is critical because not every location has a supercomputer, and given the limited amount of time, it becomes necessary to send large programs and databases throughout the network in and from smaller computers including workstations. Thus, it is fairly easy to predict that the future use of the research network will be different, though a vastly expanded version of the past. Simply extrapolating illustrates a major problem. John Rice describes communications as a significant barrier to the use of supercomputers (1987). Rice illustrates this by examining the growth in computer performance (Table1), showing how communications limit the amount of data that can be transferred (Table 2) and posing three problems which require varying amounts of communications (Table 3) together with the time to transfer the data for review (Table 4).

Table 1. Past and projected growth in raw supercomptuer speeds

year	Speed (MFLOPS)	10 year	20 year increase
1966	1	-	-
1975	4	4	-
1980	10	5	-
1985	100	25	100
1990	2,000	200	1,000
1995	200,000	2,000	50,000

Table 2. Peak and effective transfer rates of various facilities

Telephone	300	300
-----------	-----	-----

2400 baud line	2,400	2,400
9,600 baud line	9,600	9,600
Arpanet	56K	20K
Ethernet	10M	1.5M
...		

Table 3. Three problems requiring varying amounts of computation and output data

- Application 1. 1983. The high-speed impact of two steel cubes into a block of aluminum. The answer requires 3 gigabytes of pictorial data for each 30 minutes of computation.
- Application 2. 1985. Accretion of materia into a black hole (two-dimensional model). The answer requires 9 gigabytes of data in a color movie to show results.
- Application 3. 1995. Tank Battle simulation. About 8 terabytes of data are required.

Table 4. Transfer time for three applications

Facility	1983	1985	1995
Telephone	3 yrs	9 yrs	6 millenia
9,600 baud	1 mo	3 mos	2 centuries
Arpanet	2 wks	7 wks	1 century
Ethernet	5 hrs	15 hrs	16 mos

Run time                    30 mins        1 hr        1 hr

#### Workstations for Inter-Communication and Supercomputing

A recent report to the National Science Foundation on Visualization pointed out the extensive needs for networking in order to be able to visualize output from supercomputers and to intercommunicate with high quality graphical data including high resolution photographs. Since a high resolution workstation has approximately one million pixels which can change at rates of 10 to 60 per second, the bandwidth required for dynamic pictures varies from 10 to 60 million bits/sec for black and white pixels and 320 to 1920 million bits per second for 32-bit color pixels. Of course by encoding, the data-rates are reduced by several orders of magnitude, depending on the information.

By having several hundred, simple workstations operating at one megabits/sec connected to high performance computers would require a network capable of delivering 100's of Mbits/sec. The use of workstations for pictorial display and program development represents minimum requirements. In order to view an animated, color movie at 10 frames per second requires 320 Mbits/sec.

In order to view a volume of pixels (voxel) requires on the order of 4 Gigabytes of data (4 x 1000 x 1000 x 1000) which would be transferred among workstations or from supercomputers to workstations. Voxel data is typical of mechanical parts, medical images and geological data. A 45 megabit/sec link would allow a voxel to be transmitted in a few minutes provided data encoding is used.

Collaboration Technology is the name given to a range of applications whereby a community are able to work together in a distance independent fashion. These applications include: compound document transfer and simultaneous viewing and editing where the document is a combination

of text, graphics, pictures, and voice; computer conferencing involving multiple attendees with pictorial and voice interaction; design reviews via a common document; and the ability to control and interact with special laboratory and industrial facilities remotely (e.g. test equipment, factory processes and assembly lines, space station and other remote experimental facilities). NSF has a pilot project for transferring proposal information throughout the network for review, but the current line speed is insufficient for full-scale deployment.

By having having the rapid interaction that a high speed network would provide would completely change the nature of research. Today, researchers sharing a common Local Area Network are tightly connected to one another and have instantaneous communication. Because of the network overloading, connections to remote sites is slow and tenuous at best. Today, mail is the only form of communication, and in extreme cases of overload the U.S. Mail and certainly Express Mail is preferred to computer mail. If the network is used for face-to-face communication using highly encoded video data to reduce the bandwidth to 50 Kbits/sec, we would expect 100's of interations at each of the site, or an aggregate switching demand of 5 Gigabits/sec.

Finally, we expect a number of non-expected results to begin to occur when the Phase II network is installed. Campus Area Networks operating at these speeds today should provide a leading edge indication of such applications. New applications could take any form from a method to share courses and even lectures, to utilizing a collection of machines on a common problem including real time simulation.

If we use the original ARPAnet as a predictor of future, it is safe to assume that the National Research Network is likely to have both more direct impact on the research and education community, and indirect impact through the construction of a modern communications network than any other single program that can be identified.

## International Efforts In High Speed Computer Networks

Both Japan and Europe are actively building high speed (fiber optic) computer networks aimed at gigabit data transmission. In addition to the extensive projects in each country, the research program of ESPRIT and Eureka have significant research programs at all levels including fiber optics, switching, standards, and applications.

Japan's major research networks are connected via high speed communications links, including links for the archiving and transmission of international scientific and engineering material.

Today, nearly all of U.S. regional operating companies are working with foreign equipment suppliers in joint ventures to design and build the next network, ISDN. Note the technology transfer inherent in the following user/supplier relationships:

- Pacific Northwest Bell and *Northern Telecom*
- Mountain Bell and AT&T, GTE, *Ericsson, NEC and Northern Telecom*
- Wisconsin Bell and *Siemens*
- Illinois Bell and AT&T and *Fujitsu*
- NYNEX and *Siemens and Ericsson*
- Southern Bell and AT&T, *Northern Telecom and Siemens*

While it is clear the services provided by ISDN in the next 5-10 years will be unsuitable for meeting the needs of the National Research Network,

by the action of the suppliers, the communications equipment market will become controlled by foreign supply.

### **Importance of the National Research Network for Industry**

While virtually every program emanating from Washington carries the competitiveness banner, the National Research Network is one which, based on the experience of ARPAnet, is likely to have more effect than all others including superconductors (at least for another decade). Two kinds of benefits should result: direct -- through use and new services, and indirect -- through the design collaboration among academe and the computer and communications industries to invent the communication and computer system and services for the 21st century.

By having much higher bandwidth, more direct links can be established between industry and academe. Like ARPAnet, the network will serve as the leading edge, large scale social experiment in very high bandwidth intercommunication. It will be possible to perform conduct collaborative tele-science via the network in a distributed and highly democratic fashion. The network should permit much of what we do to be space-independent. Rather than limiting the communication via low data-rate mail facilities, the National Research Network's facilities will allow "presence" and coupling between the organizations for collaborative research, common research facilities, and education. Two-way exchange and high bandwidth coupling are essential for technology transfer, where today the transfer of people is the best and really only significant method.

While ARPAnet created a new form of communication in spite of and around the communications industry, the National Research Network will involve both the computer and communications industries for standards and equipment development. Computer users require higher bandwidth in order to interconnect their Campus and Local Area Networks, and a synergy is possible and necessary based on need. Unlike the computer industry which has a history of collaboration, the telecommunications industry constituents have no collaborative research either with universities

or with consortia (e.g. MCC, SRC, or Semitech) outside of Bellcore. In order to be competitive with international tele-communications efforts, both the telecommunications and computer industries must learn to collaborate with academe, where basic research is done.

In a recent example of cooperation with industry, NSF received numerous, high quality proposals to build and operate the high speed backbone network consisting of about 20 nodes connected via T1 links. We know of at least four methods by which a high speed switch capable of operating at DS3 (45 megabits/second) can be built and installed within three years in order to attain the aggressive timetable posited in the report.

### **Current Impediments to Building a National Research Network**

Building the National Research Network is a very slow process, given its history, and given that the ARPAnet has operated at essentially constant speed for nearly 15 years, independent of the radical change in the nature of the computing nodes that it interconnects.

In addition to the agency inertia which comes with the size of the problem being addressed, a number of other factors mitigate against progress.

These argue even more for the government to take a leadership role, in what we believe to be a crucial national facility.

#### **Communications Carriers and Equipment Suppliers**

The suppliers of communications services have a large built-in market and an installed network which creates inertia based on voice traffic.

Carriers blame regulatory agencies for the voice-based pricing which virtually requires any organization wishing to communicate, to establish a private network for data traffic. Altogether, we have lack of a perceived need for a high speed network capable of interconnecting Local and Campus Area Networks, especially aimed at research. Fortunately, the rapid growth in LANs in the commercial segment is beginning to create a market for DS1 (1.5 Mbits) and DS2 (6 Mbits) links and switched links.



Unfortunately, the carriers believe they have a solution in the up and coming ISDN. Virtually no one is complimentary about ISDN either in terms of the schedule of availability or that it is well enough defined among all local and international carriers. There is significant competition among AT&T and the Regional Operating Companies to set proprietary standards to enable monopolies. Also when high speed fibers do terminate in the switching offices, it is difficult to get the fibers distributed to the campuses because of the famous "the last mile problem". For the National Research Network, ISDN is generally irrelevant. It is a factor of 200-2000 slower than the local and campus area networks and a factor of 1000 slower than DS3 (45 Mbits/sec) and a factor of 30,000 slower than is potentially viable for fiber. We require DS3 link capacity today at every campus.

Given the lack of interest by the carriers for providing high speed switching services at T1 and above speeds, the equipment suppliers simply aren't providing equipment. Similarly, proposals for standards for speeds higher than T1 are moving at sub-glacial rates. (Glacial is the standard unit of measure of motion for government action.) Only recently, have T1 packet switches become available.

The situation can simply be characterized as a traditional "technology push versus market pull" problem -- new, high speed (high technology) services are not available because we have no demand... and we have no demand because the services are not available.

Inherent National Inabilities to Fund A National Research Network

Unlike the Department of Transportation which has responsibility for our highway network, no federal agency has responsibility for supplying networking services. Also, given the rapid changes in technology, such a network would most likely be perennially obsolete (using the FTS Voice

and Data Communications Networks as a guide), and hence require continued funding.

While the Subcommittee recommended a lead agency to secure and manage the network, it is unclear that any agency would take on this responsibility since a network will cut into the budget funding for regular mission activities, including research within the agency in a zero sum fashion, even though it is a facility for the entire research establishment. We have no examples of a single agency supplying a facility for the entire research enterprise.... the behaviour is byzantine (each agency wants its own facility, no matter what the cost). The Federal organization and budget is simply incapable of setting up a line item to be funded and administered either by an agency or an interagency group such as the FCCSET subcommittee as a facility for all of the research and higher education community.

Ideally, a single line item for a number of common facilities could be separated out and managed for all agencies. Supercomputers and the National Research Network should both be funded and managed in this fashion. By not having a common facilities budget for all of science, each agency is obliged to fund and build its own, overlapping networks and computer centers. Today, each agency has roughly two networks for intercommunication among the research community. Similarly, each agency has or is in the process of constructing agency-specific supercomputer centers.

In addition to agency specific networks and facilities, divisions, communities of interest have and operate specific networks in an overlapping, non-communicating fashion. The situation is often likened to telephone networks prior to Vail when a town might have several separate, non-communicating networks requiring a subscriber to have several phones on a desk. Today, a campus such as MIT has 10 connections to separate federal networks.

### Competition for Other Research Funds

The National Research Network competes or rather should compete for funds with other facets of research, and research facilities. The supercomputer centers invariably question the need for a high performance network, but ironically, the concept of a few, regional supercomputer centers is useless without a modern high speed network, to transmit pictures and files, which is substantially better than we have today. It is necessary to couple workstations to supercomputers for a truly productive and cost-effective supercomputing facility.

On the other hand, advocates of smaller computers, including mini-supercomputers, would simply eliminate the centers and hence the need for the network. It is critical to have a heirarchy of computers so that a user can compute at the site that meet the requirements.

### **Moving Ahead Toward The National Research Network**

The report, recommending a National Research Initiative to OSTP was not forwarded to congress, but instead OSTP asked Congress for an extension in order to incorporate the proposal for a National Research Network into a much larger report encompassing the support of computer science and the support of an entire infrastructure for computer applications in science and engineering. A new report was initiated, based on the FCCSET Committee on Computer Research and Applications subcommittees on: Science and Engineering Computing; Computer Research and Development; and Computer Networking, Infrastructure, and Digital Communications. The report from the FCCSET Committee on Computer Research and Applications is the National Computing Strategy (NCS), using as a model, the National Aeronautical R&D Goals which outlined a series of eight actions to maintain U.S. leadership in aeronautics.

Given the unclear record of very large programs (e.g. Space Station, Star Wars, Superconducting Super-Collider, Superconductors, Human Genome Mapping) in achieving funding, managing, and achieving meaningful results from such programs, once the funds are secured, it is important to understand the motivation of the National Research Network Initiative so that it can proceed as rapidly as possible, independent of the success or failure of a much larger set of goals as incorporated in a broad report like the National Computing Strategy.

Given the inability to attain glacial speed, we simply must not lose yet another year through inaction. The report outlined a single method, selecting a lead agency and allocating the budget to this agency, as the method to build and operate the Network. The perils of this kind of funding and operation are numerous and constitute a major impediment as previously discussed.

We believe the National Research Network can and must proceed using the current interagency coordinating committee. Even without the impetus of a major government initiative and outlay, funds can and must be made available on a shared basis utilizing existing agency, industrial, and academic resources in a much more propitious fashion. It is highly possible that simply coordinating in this fashion will save money for all agencies by not having to operate duplicate networks. Hence, the resolve for a network through the coordinated effort of the agencies is paramount.

# A National Research Network: Round I

Gordon Bell

## Introduction

Modern science depends on rapid communications and information exchange. Today, some thirty-six major national and international networks exist using either the ARPAnet packet switching ideas, or a modified form of packet switching which is carried out directly by each host computer. State and regional networks are proliferating. NSFNET, an "internet" designed initially to improve access to supercomputer centers, has in the space of two years, forged links among 17 state, regional, and federal agency networks.

In the early '80's the lack of access to supercomputing power by the scientific and engineering community caused the formation of the NSF Office of Advanced Scientific Computing, which funded five centers for supercomputers. Given the highly distributed location of users, the need for a national wide area network for computer access and for the interchange of associated scientific information (mail, files, databases, etc.) became clear. It immediately became obvious that existing agency networks both lacked the inherent capacity and were overloaded. In fact, the current Federal wide area networks still operate at the speed (56 Kbits/sec) of the original ARPAnet (circa 1972), despite the fact that computers have both increased in number (by a factor of 100) and speed (by a factor of 30). Moreover, the very nature of what is communicated has changed with the power of computers. Computer simulations of physical process yield output expressed as three dimensional dynamic process. Computer-based research tools are becoming more interactive, and the problems under study often require the active collaboration of researchers who are distributed in various research institutions. Network technology has fallen seriously behind computing innovations and the subsequent requirements of the research enterprise.

Today's fiber optic communications links offer the ability for a factor of 1000-100,000 increase in capacity

and speed over traditional cable and satellite channels. The price of fiber optic links are based on equivalent voice grade circuits, and as such remain high, despite a vast amount of unused fiber. Furthermore, switching equipment to exploit even DS3 (45 Mbits/sec) is unavailable. While several proposals for DS3 fast packet protocols and switches exist, no technology exists to switch packets at fiber optic speeds. Research and development work must be expanded for fiber optic switching.

Most Federal agencies support an average of two, independent wide area networks that couple researchers to one another for mail, collaboration, and file transfer and to large, central systems. Often these networks go to the same campus (an academic institution, federal or industrial lab), but to a different building, wasting resources in the process. While a few years ago, all wide area networks used different protocols and could not communicate with one another, today most networks are migrating toward DARPA's TCP/IP protocol with a commitment to use the ISO protocols when available. By having identical protocols, networks can be interconnected together easily (internetting) and share common links to save costs.

Campuses are being wired to interconnect the array of Local Area Networks- LANs (a wiring scheme for a single building or small cluster of buildings) to form Campus Area Networks - CANs. Since LANs operate at 10 Mbits/sec and CANs operate at up to 100 Mbits/sec. a very fast wide area network is required in order to interconnect the CANs (and LANs) to form what will become a Global, Local Area Network. At present, these Campus Area Networks form networked islands. They require a much higher degree of wide and global area networking in order to communicate with other institutions, i.e. a National Research Network.

### **Congressional Action and Executive Branch Response**

Some in Congress have recognized the problem and the opportunity. In June 1986, Senator Gore introduced the Supercomputer Network Study Act of 1986. Public Law 99-383, passed June 21, 1986 by the 99th Congress, charged the Office of Science and Technology Policy (OSTP) with conducting "a

study of critical problems and current and future options regarding communications networks for research computers, including supercomputers at universities and federal research facilities in the United States".

At OSTP's direction, an interagency group under the auspices of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) for Computer Research and Applications formed to carry out the study of the following issues:

- the networking needs of the nation's academic and federal research computer programs, including supercomputer programs, over the next 15 years, including requirements in terms of volume of data, reliability of transmission, software compatibility, graphics capabilities, and transmission security;
- the benefits and opportunities that an improved computer network would offer for electronic mail, file transfer, and remote access and communications; and
- the networking options available for linking academic and research computer, including supercomputers, with a particular emphasis on fiber optics.

The Computer Network Study Planning Group, composed of participants from DoD (including DARPA), DoE, HHS, NASA, NBS, and NSF produced and delivered a three volume report for OSTP in early August 1987, consisting of the following documents:

- Volume I. Recommendations to Congress for the National Research Network
- Volume II. A workshop involving 100 researchers, network users, suppliers, and policy officials on the network giving requirements and alternatives, special requirements for supercomputers, internet concepts, future standards and service requirements, computer security, and the government role.
- Volume III. Background papers on internet concepts, network requirements by various agencies, and government's role

### **Summary of Recommendations to OSTP (for Congress) for the National Research Network**

The Subcommittee on Networking, Infrastructure and Digital Communications (NIDC) recommended the following:

"The U.S. should undertake, as a national goal, the establishment of a National Research Network in a staged approach that supports the upgrade of current facilities, and development of needed new capabilities. Achievement of this goal would foster and enhance the U.S. position of world leadership in computer networking.

As rapidly as feasible, the National Research Network should be designed, deployed and maintained as an advanced computer network. This network should interconnect substantially every academic, industrial, and government research establishment and unique scientific resource to encourage scientific collaboration unhindered by distance and to permit the sharing of unique research facilities and resources. Since security of the network is a vital concern, appropriate policies should be adopted to protect the information in the network from threats, vulnerabilities and risks, and to assure a uniform level of security.

Until the National Research Network can replace the current system, existing networks should be maintained and modified as they join the national network. Since supercomputer systems comprise a special and valuable national research resource with very high performance requirements, the responsibility for network access to supercomputers should be vested in the supercomputer centers themselves until the advanced computer network, capable of offering the requisite service level, is operational.

Industry should be encouraged through special incentives to participate in research, development, and deployment of the National Research Network. Tariff schedules which have been set for voice transmission should be re-examined in light of the requirements for transmission of data through computer networking.



To meet the goal for the National Research Network and to set and agenda for the future the following actions are recommended:

- The Subcommittee on Computer Networking, Infrastructure and Digital Communications which was established by the Office of Science and Technology Policy on May 15, 1987, should oversee the first stage in development and operation of the National Research Network -- a coordinated internetwork that would include the Federal agencies that currently operate research supporting networking.
- The FCCSET Subcommittee on Computer Networking, Infrastructure and Digital Communications should identify a lead agency which would be responsible for requesting funds for the National Research Network, and eventually for selecting a contractor to manage the network. The manager would be responsible for Stages 2 (45 Mbit/sec links) and Stage 3 (multi-Gigbit/sec links).
- As a first stage in the development of the National Research Network, the current Internet system developed by DARPA and networks supported by agencies should be interconnected over the next two years. These facilities, if coordinated and centrally managed, have the capability to interconnect many computer networks into a single virtual computer network. The Federal government should encourage and assist research facilities and academic institutions to establish local and campus area networks to connect to the Internet systems. The estimated cost for this proposed upgraded service is \$5 million per year, and should be implemented through the shared resources of NSF, DoE, DARPA, NASA, and HHS.
- In the second stage, new funding for development should be requested at \$5 million per year over the next five years to upgrade and expand the nation's existing computer networks, which support research programs, to achieve data communications at 1.5 Mbits/sec to 200-400 U.S. research institutions. It is estimated that these expanded and upgraded facilities will require an additional funding of approximately \$50,000,000 to operate (GB: assumed the price of T1 lines

would decline by a factor of 2 over the next five years, whereas, in the six months since the report, line charges for many T1 lines have already dropped more than a factor of two due to oversupply.)

- In the third stage, a vigorous focused program of research and development for the National Research Network should be immediately established. A total of \$400 million is needed over ten years to advance the knowledge base and technology of computer network capabilities in order to achieve data communications and switching capabilities to support transmission of three billion bits per second within fifteen years. This will support a network 100,000 times more capable than currently available and will be essential to foster scientific collaboration and sharing of research resources. When fully deployed, the cost of operating this advanced network is estimated to be \$400 million per year, given the current commercial tariffs for data communications.

Support should be given to the development of standards and their harmonization in the international arena. Aggressive action is needed to increase user participation in the standards development process, to get requirements for standards expressed early in the development process, and to speed the implementation of standards in commercial off-the-shelf products. It is essential that standards development be carried out within the framework of overall systems requirements to achieve interoperability, common user interfaces to systems and enhanced security."

The structure of the Network as a hierarchy of networks forming an Internet, and the timetable for achieving the first implementations at a given speed is shown in Figures 1 and 2 from the Report.

### **Motivation for the National Research Network ... The Uses**

If we use the original ARPANet as a predictor of future, it is safe to assume that the National Research Network is likely to have both more direct impact on the research and education community, and indirect impact through the construction of a modern communications network than any other single program that

can be identified.

ARPAnet was initially justified on the basis of being able to share facilities including particular programs, databases, and files, and being able to promote the use of idle computer capacity. What happened was that a completely different style of interaction developed based on being able to send mail and large documents. Extensive public forums and conferences evolved through bulletin board and computer conferencing. Remote terminal access is negligible, and file interchange was relatively smaller than expected.

With supercomputing, remote access is critical. Few institutions can support supercomputers, and given the limited amount of time, it becomes necessary to send large programs and databases throughout the network in and from smaller computers including workstations. Thus, it is fairly easy to predict that the future use of the research network will be different, though a vastly expanded version of the past. Simply extrapolating from some new applications illustrates the problem.

A recent report to the National Science Foundation on Visualization pointed out the extensive needs for networking in order to be able to visualize output from supercomputers and to intercommunicate with high quality graphical data including high resolution photographs. Since a high resolution workstation has approximately one million pixels which can change at rates of 10 to 60 per second, the bandwidth required for dynamic pictures varies from 10 to 60 million bits/sec for black and white pixels and 320 to 1920 million bits per second for 32-bit color pixels. Of course by encoding, the data-rates are reduced by several orders of magnitude, depending on the information.

By having several hundred, simple workstations operating at one megabits/sec connected to high performance computers would require a network capable of delivering 100's of Mbits/sec. The use of workstations for pictorial display and program development represents minimum requirements. In order to view an animated, color movie at 10 frames per second requires 320 Mbits/sec. In order to view a

volume of pixels (voxel) requires on the order of 4 Gigabytes of data (4 x 1000 x 1000 x 1000) which would be transferred among workstations or from supercomputers to workstations. Voxel data is typical of mechanical parts, medical images and geological data. A 45 megabit/sec link would allow a voxel to be transmitted in a few minutes provided data encoding is used.

Collaboration Technology is the name given to a range of applications whereby members of a community are able to work together in a distance independent fashion. These applications include: compound document transfer and simultaneous viewing and editing where the document is a combination of text, graphics, pictures, and voice; computer conferencing involving multiple attendees with pictorial and voice interaction; design reviews via a common document; and the ability to control and interact with special laboratory and industrial facilities remotely (e.g. test equipment, factory processes and assembly lines, space station and other remote experimental facilities). NSF has a pilot project for transferring proposal information throughout the network for review, but the current line speed is insufficient for full-scale deployment.

By having having the rapid interaction that a high speed network would provide would completely change the nature of research. Today, researchers sharing a common Local Area Network are tightly connected to one another and have instantaneous communication. Because of the network overloading, connections to remote sites is slow and tenuous at best. In extreme cases of overload the U.S. Mail and certainly Express Mail is preferred to computer mail. If the network is used for face-to-face communication using highly encoded video data to reduce the bandwidth to 50 Kbits/sec, we would expect 100's of interations at each of the site, or an aggregate switching demand of 5 Gigabits/sec.

Finally, we expect a number of unanticipated applications to begin to develop when the Phase II network is installed. Campus Area Networks operating at these speeds today should provide a leading edge indication of such applications. Even in Phase I, with full interconnectivity at T1 speeds for the backbone and major regional networks, a new set of library, database, common courseware and educational

services will be possible. New applications could take any form from a method to share courses and even lectures, to utilizing an experimental facility or even a collection of machines on a common problem including real time simulation for both research and educational purposes.

### **Industrial Motivation**

The interconnection among government laboratories, universities, and industrial research organizations provided by the Network should facilitate technology transfer unlike almost any other facility. It will enable the instantaneous transfer of much of our knowledge which we are increasingly embedding in databases and programs. Also, all of the other uses described above will be available throughout the community.

In addition to the direct benefit resulting in the use of the network, we believe the network is vital for maintaining a competitive telecommunications industry. By interacting with the academic community to build and evolve the network, we expect to build the same vitality and synergy that the computer industry enjoys.

Both Japan and Europe are actively building high speed (fiber optic) computer networks aimed at gigabit data transmission. In addition to the extensive projects in each country, the European research program of ESPRIT and Eureka have significant research programs at all levels including fiber optics, switching, standards, and applications.

Japan's major research networks are connected via high speed communications links, including links for the archiving and transmission of international scientific and engineering material.

Today, nearly all of U.S. regional operating companies are working with foreign equipment suppliers in joint ventures to design and build the next network, ISDN. Note the technology transfer inherent in the following user/supplier relationships:

- Pacific Northwest Bell and *Northern Telecom*
- Mountain Bell and AT&T, GTE, *Ericsson, NEC and Northern Telecom*
- Wisconsin Bell and *Siemens*
- Illinois Bell and AT&T and *Fujitsu*

- NYNEX and Siemens and *Ericsson*
- Southern Bell and AT&T, *Northern Telecom and Siemens*

While it is clear the services provided by ISDN in the next 5-10 years will be unsuitable for meeting the needs of the National Research Network, by the action of the suppliers, the communications equipment market will become controlled by foreign supply.

The situation today in computer networking is a classical technology push, market pull dilemma. New, high speed (high technology) services are not available because we have no demand, and we have no demand because the services are not available.

By not having public, high capacity network switching services available, every company, company division, laboratory, agency, division of an agency, etc. must acquire lines, acquire switching equipment, install, and operate its own data network. The situation is not unlike the telephone network prior to Vail, when a town might have several non-communicating company networks requiring a subscriber to have several phones on a desk. Today, a large research campus has over a dozen, incompatible networks of essentially the same topology.

### **Moving Ahead Toward The National Research Network**

To date, the report, recommending a National Research Initiative to OSTP was not forwarded to Congress. Instead OSTP has asked Congress for an extension in order to incorporate the proposal for a National Research Network into a much larger report encompassing the support of high performance computing, computer science, and revamped infrastructure for computer applications in science and engineering. This new report from the FCCSET Committee on Computer Research and Applications was based on an integration of its subcommittee reports: Science and Engineering Computing; Computer Research and Development; and Computer Networking, Infrastructure, and Digital Communications has been submitted to OSTP. The report, "A National R&D Strategy for High Performance Computing".

Given the clouded record of very large programs (e.g. Space Station, Star Wars, Superconducting Super-Collider, Superconductors, Human Genome Mapping) in either achieving funding, or achieving meaningful results, one can hardly be optimistic about the outcome of a large scale recommendation encompassing virtually every facet of computer science, technology, use, and networking.

The report outlined a single method, selecting a lead agency and allocating the budget to this agency, as the method to build and operate the Network. Again, the perils of this kind of funding and operation are numerous and constitute another impediment since in an environment of zero sum agency budgeting, the responsibility for a new network will not be accompanied by a budget, but rather the budget will come from the agency's current commitments (e.g. research, facilities).

The scientific community (the research and education community users and network providers) must understand the motivation of the National Research Network initiative so work can proceed as rapidly as possible, independent of the success or failure of a much larger set of goals for higher performance computing. We simply must not lose yet another year through inaction.

Unlike the superhighway system, any one or a combination of the existing telecommunications suppliers could pre-empt the National Research Network activity and simply build the network and offer the service for sale. A highly aggressive, imaginative industry could view the Network as the major, large scale, social experiment of the century in preparation for the 21st century. An effort to achieve Phase II (45 Mbits/sec) in three years would be a small fraction of the expense of this industry's research, development, and operations budget. Government should encourage efforts of this kind in every way possible, but, based on history, we certainly can't depend on it.

We believe the National Research Network can and must proceed using the current interagency coordinating committee. Even without the impetus of a major government initiative and outlay, funds can and must be made available on a shared basis utilizing existing agency, industrial, and academic

resources in a much more propitious fashion. It is highly possible that simply coordinating in this fashion will save money for all agencies by not having to operate duplicate networks. Hence, the resolve for a network through the coordinated effort of the agencies is paramount.



## Steps Toward A National Research Network

Gordon Bell

### Dana Computer Sunnyvale, California

#### Introduction

Modern science depends on rapid communications and information exchange. Today, many major national and international networks exist using some form of packet switching to interconnect host computers. State and regional networks are proliferating. NSFNET, an "internet" designed initially to improve access to supercomputer centers, has in the space of two years, forged links among 17 state, regional, and federal agency networks.

In the early '80's the lack of access to supercomputing power by the research community caused the formation of the NSF Office of Advanced Scientific Computing, which funded five centers for supercomputers. Given the highly distributed location of users, the need for a national wide area network for computer access and for the interchange of associated scientific information (mail, files, databases, etc.) became clear.

It immediately became obvious that existing agency networks both lacked the inherent capacity and were overloaded. In fact, the current Federal wide area networks still operate at the speed (56 Kbits/sec) of the original ARPAnet (circa 1972), despite the fact that computers have both increased in number (by a factor of 100) and speed (by a factor of 30). Moreover, the very nature of what is communicated has changed with the power of computers. Computer simulations of physical processes yield output expressed as three dimensional dynamic graphics. Computer-based research tools are becoming more interactive, and the problems under study often require the active collaboration of researchers who are distributed in various research institutions.

Today's fiber optic communications links offer the ability for a factor of 1000-100,000 increase in capacity and speed over traditional cable and satellite channels. The price of fiber optic links are based on equivalent voice grade circuits, and as such remain high, despite a vast amount of unused fiber. Switching equipment to exploit even DS3 (45 Mbits/sec) is unavailable. While several proposals for DS3 fast packet protocols and switches exist, no technology exists to switch packets at fiber optic speeds.

Most Federal agencies support an average of two, independent wide area networks that couple researchers to one another for mail, collaboration, file transfer and to large, central systems. Often these networks go to the same campus (an academic institution, federal or industrial lab), but to a different building, wasting resources in the process. While a few years ago, all wide area networks used different protocols and could not communicate with one another, today most networks are migrating toward DARPA's TCP/IP protocol with a commitment to use the ISO protocols when available.

Campuses are being wired to interconnect the array of Local Area Networks- LANs (a wiring scheme for a single building or small cluster of buildings) to form Campus Area Networks - CANs. At present, these Campus Area Networks form networked islands.

#### Congressional Action and Executive Branch Response

Some in Congress have recognized the problem and the opportunity. In June 1986, Senator Gore introduced the Supercomputer Network Study Act of 1986. Public Law 99-383, passed June 21, 1986 by the 99th Congress, charged the Office of Science and Technology Policy (OSTP) with conducting "a study of critical problems and current and future options regarding communications networks for research computers, including supercomputers at universities and federal research facilities in the United States".

At OSTP's direction, an inter-agency group under the auspices of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET) for Computer Research and Applications formed to carry out the study of the following issues:

- the networking needs of the nation's academic and federal research computer programs, including supercomputer programs, over the next 15 years, including requirements in terms of volume of data, reliability of transmission, software compatibility, graphics capabilities, and transmission security;
- the benefits and opportunities that an improved computer network would offer for electronic mail, file transfer, and remote access and communications; and
- the networking options available for linking academic and research computer, including supercomputers, with a particular emphasis on fiber optics.

The Computer Network Study Planning Group, composed of participants from DoD (including DARPA), DoE, HHS, NASA, NBS, and NSF produced and delivered a three volume report for OSTP in early August 1987, consisting of recommendations, result of a workshop, and background papers.

#### **Summary of Recommendations to OSTP (for Congress) for the National Research Network**

The Subcommittee on Networking, Infrastructure and Digital Communications (NIDC) recommended the following:

"The U.S. should undertake, as a national goal, the establishment of a National Research Network in a staged approach that supports the upgrade of current facilities, and development of needed new capabilities. Achievement of this goal would foster and enhance the U.S. position of world leadership in computer networking.

As rapidly as feasible, the National Research Network should be designed, deployed and maintained as an advanced computer network. This network should interconnect substantially every academic, industrial, and government research establishment and unique scientific resource to encourage scientific collaboration unhindered by distance and to permit the sharing of unquiet research facilities and resources. Since security of the network is a vital concern, appropriate policies should be adopted to protect the information in the network from threats, vulnerabilities and risks, and to assure a uniform level of security.

Until the National Research Network can replace the current system, existing networks should be maintained and modified as they join the national network. Since supercomputer systems comprise a special and valuable national research resource with very high performance requirements, the responsibility for network access to supercomputers should be vested in the supercomputer centers themselves until the advanced computer network, capable of offering the requisite service level, is operational.

Industry should be encouraged through special incentives to participate in research, development, and deployment of the National Research Network. Tariff schedules which have been set for voice transmission should be re-examined in light of the requirements for transmission of data through computer networking.

To meet the goal for the National Research Network and to set an agenda for the future the following actions are recommended:

- The Subcommittee on Computer Networking, Infrastructure and Digital Communications which was established by the Office of Science and Technology Policy on May 15, 1987, should oversee the first stage in development and operation of the National Research Network -- a coordinated internetwork that would include the Federal agencies that currently operate research supporting networking.
- The FCCSET Subcommittee on Computer Networking, Infrastructure and Digital Communications should identify a lead agency which would be responsible for requesting funds for the National Research Network, and eventually for selecting a contractor to manage the network. The manager would be responsible for Stages 2 (45 Mbit/sec links) and Stage 3 (multi-Gigabit/sec links).
- As a first stage in the development of the National Research Network, the current Internet system developed by DARPA and networks supported by agencies should be interconnected over the next two years. These facilities, if coordinated and centrally managed, have the capability to interconnect many computer networks into a single virtual computer network. The Federal government should encourage and assist research facilities and academic institutions to establish local and campus area networks to connect to the Internet systems. The estimated cost for this proposed upgraded service is \$5 million per year, and should be implemented through the shared resources of NSF, DoE, DARPA, NASA, and HHS.
- In the second stage, new funding for development should be requested at \$5 million per year over the next five years to upgrade and expand the nation's existing computer networks, which support research programs, to achieve data communications at 1.5 Mbits/sec to 200-400 U.S. research institutions. It is estimated that these expanded and upgraded facilities will require an additional funding of approximately \$50,000,000 to operate (GB: assumed the price of T1 lines would decline by a factor of 2 over the next five years, whereas, in the six months since the report, line charges for many T1 lines have already dropped more than a factor of two due to over-supply.)
- In the third stage, a vigorous focused program of research and development for the National Research Network should be immediately

established. A total of \$400 million is needed over ten years to advance the knowledge base and technology of computer network capabilities in order to achieve data communications and switching capabilities to support transmission of three billion bits per second within fifteen years. This will support a network 100,000 times more capable than currently available and will be essential to foster scientific collaboration and sharing of research resources. When fully deployed, the cost of operating this advanced network is estimated to be \$400 million per year, given the current commercial tariffs for data communications.

Support should be given to the development of standards and their harmonization in the international arena. Aggressive action is needed to increase user participation in the standards development process, to get requirements for standards expressed early in the development process, and to speed the implementation of standards in commercial off-the-shelf products. It is essential that standards development be carried out within the framework of overall systems requirements to achieve interoperability, common user interfaces to systems and enhanced security."

### **Motivation for the National Research Network**

If we use the original ARPAnet as a predictor of future, it is safe to assume that the National Research Network is likely to have both more direct impact on the research and education community, and indirect impact through the construction of a modern communications network than almost any other single program that can be identified.

ARPAnet was initially justified on the basis of being able to share facilities including particular programs, databases, and files, and being able to promote the use of idle computer capacity. What happened was that a completely different style of interaction developed based on being able to send mail and large documents. Extensive public forums and conferences evolved through bulletin board and computer conferencing. Remote terminal access is negligible, and file interchange was relatively smaller than expected.

With supercomputing, remote access is critical. Few institutions can support supercomputers, and given the limited amount of time, it becomes necessary to send large programs and databases throughout the network in and from smaller computers including workstations. A recent NSF report on Visualization pointed out the need for research which would enable users to deal with the prodigious amount of data that comes from modern computers and applications. Also, we would expect a new form of video and computer conferencing to be possible utilizing high speed networks. Already researchers have defined this to be collaboration technology. Of course, a revolution is possible in how libraries might share information. Thus, it is fairly easy to predict that the future use of the research network will be different, though a vastly expanded version of the past.

The inter-connection among government laboratories, universities, and industrial research organizations provided by the Network should facilitate technology transfer unlike almost any other facility. It will enable the instantaneous transfer of much of our knowledge which we are increasingly embedding in databases and programs. Unlike earlier networks, the National Research Network is aimed at connections with many more institutions.

In addition to the direct benefit resulting in the use of the network, we believe the network is vital for maintaining a competitive telecommunications industry. By interacting with the academic community to build and evolve the network, we expect to build the same vitality and synergy that the computer industry enjoys.

Other countries are actively building high speed (fiber optic) computer networks aimed at gigabit data transmission. For example, the European research program of ESPRIT and Eureka have significant research programs at all levels including fiber optics, switching, standards, and applications.

### **Onward Toward The National Research Network**

OSTP asked Congress for a time extension until November in order to incorporate the proposal for a National Research Network into a much larger report encompassing the support of high performance computing, computer science, and revamped infrastructure for computer applications in science and engineering. This new report from the FCCSET Committee on Computer Research and Applications was based on an integration of its subcommittee reports: Science and Engineering Computing; Computer Research and Development; and Computer Networking, Infrastructure, and Digital Communications has been submitted to OSTP. The resulting report is entitled "A National R&D Strategy for High Performance Computing".

Given the clouded record of very large programs (e.g. Space Station, Star Wars, Superconducting Super-Collider, Superconductors, Human Genome Mapping) in either achieving funding, or achieving meaningful results, one can hardly be optimistic about the outcome of a large scale recommendations encompassing virtually every facet of computer science, technology, use, and networking.

The community (the research and education community users and network providers) must understand the motivation of the National Research Network initiative so work can proceed as rapidly as possible, independent of the success or failure of a much larger set of goals for higher performance computing. We simply must not lose yet another year through inaction or in coupling with a larger program.

Unlike the super-highway system, any one or a combination of the existing telecommunications suppliers could pre-empt the National Research Network activity and simply build the network and offer the service for sale. A highly aggressive, imaginative industry could view the Network as the major, large scale, social experiment of the century in preparation for the 21st century. An effort to achieve Phase II (45 Mbits/sec) in three years would be a small fraction of the expense of this industry's research, development, and operations budget. Government should encourage efforts of this kind in every way possible, but, based on history, we certainly can't depend on **industry**. NSF has just established the first part of what could be the network in its backbone network.

The National Research Network can and must proceed using the current inter-agency coordinating committee. Even without the impetus of a major government initiative and outlay, funds can and must be made available on a shared basis utilizing existing agency, industrial, and academic resources in a much more propitious fashion. It is highly possible that simply coordinating in this fashion will save money for all agencies by not having to operate duplicate networks. Hence, the resolve for a network through the coordinated effort of the agencies is paramount. Now is a great time for government to take a first step in considering its constituents and supporters (taxpayers).

**From:** Gordon Bell  
**To:** [Andreu Vea Baro \(andreu@veabaro.info\)](mailto:andreu@veabaro.info)  
**Subject:** FW: Questions about NSF when CISE was formed  
**Date:** Sunday, November 7, 2004 10:18:00 AM

---

Note this correspondence.

g

---

**From:** Gordon Bell  
**Sent:** Monday, July 07, 1997 4:32 PM  
**To:** 'Jed Gordon'  
**Cc:** 'Charles Brownstein'  
**Subject:** RE: Questions about NSF when CISE was formed

I hope you'll talk to Erich Bloch and Chuck on this. Also, Bernie Chern and other folks who were in the founding like Mel Ciment.

-----Original Message-----

**From:** Jed Gordon [[SMTP:JGordon@nas.edu](mailto:JGordon@nas.edu)]  
**Sent:** Monday, July 07, 1997 1:44 PM  
**To:** Gordon Bell  
**Cc:** Jerry Sheehan  
**Subject:** RE: Questions about NSF when CISE was formed

What events precipitated the establishment of CISE?

**[Gordon Bell]** Eric Bloch had come to NSF and wanted to strengthen computer science, together with Engineering to have a computing directorate. The supercomputing initiative was just starting up where several centers had been funded to supply supers to scientists. Computing was in all of the directorates and the computing community wanted to have this brought together. Recall Computer Science had the largest growth in departmental size in universities when CISE was formed. Also, we were not producing enough PhDs to keep up with either industry or educational needs.

I attach a fax of an interview in IEEE Software, July 1987 when CISE started. It's in .tif format and if you need it faxed to you, then send me a fax number.

How did CISE differ from its predecessors?

**[Gordon Bell]** CISE came from Math and Physical Science where Kent Curtis headed traditional CS research. This included some work in AI. It was difficult to get funding and program attention because it was down a couple of levels in the organization. It also came from a part of Engineering where architectural research and chips were being built, headed by Bernie Chern. Robotics also was done in engineering along with some efforts in AI. There were programs in library science that came over too. Two of the largest efforts was the supercomputing centers program and its networking component.

What was your position before becoming the director of CISE?

**[Gordon Bell]** I had been head of R&D at Digital and a strong voice for supporting computer science

and engineering research at NSF. I had started a company, Encore Computer, and had just moved to Silicon Valley to work with startups.

What attracted you to the position?

**[Gordon Bell]** Erich Bloch and the chance to pull all these programs together. I had been frustrated by computing's organizational position within NSF and the fact that it was distributed in all the directorates. I also wanted to get NSF focused in a more applied fashion that could work with industry and also work on parallelism. Also, I feel everyone should spend some time working for the government in some capacity.

What role, if any, did you have in developing the organizational structure of CISE?

**[Gordon Bell]** Erich gave me various pieces to put together. He also asked Nam Suh, who headed engineering and I to work together to decide what should be in CISE from engineering. Erich made the decision. I had argued that more should come into CISE that would support communications and networking, but that remained in engineering because of my network interests. I think it is clear NOW that more communication and networking should have come into computing as computer science is pretty far removed from audio, video, and voice.

What was the reaction of the rest of NSF ( the directorate of Engineering in particular) to the creation of CISE?

**[Gordon Bell]** We (Nam Suh) argued about what should be in each part. Nam wanted to keep everything as it was... I wanted most of computing sans signal processing to come to CISE.

It should be noted that John Connolly, who headed the Supercomputers program was not especially happy as he was busily getting centers funded that we couldn't ultimately fund. I argued for fewer, larger centers or centers such as NCAR that were funded by the sciences/users or that the users should pay something for supers to encourage some sort of market mechanism. I recall a particular contentious meeting in Cornell with all the supercomputer directors present when I told them they all had to run UNIX and that they had no responsibility for networking... and they had to get more outside funding!

I cut off several of the centers that were temporary and stopped the funding of a large consortium center at Princeton because the computer they selected didn't work. They were pissed and I got all sorts of calls from congress persons and college presidents and Nobel prize winners when I pushed the idea of pay for performance. Erich supported me.

The big changes were that I took all the networking from the supers program and made it a separate division that created NSFnet, nii, nren, etc. Again the supers folks were pissed.

How did the relationship between the previous computer science programs and the electrical engineering/computing programs change?

**[Gordon Bell]** I don't think it changed much. Ask Bernie Chern on this one.

How much, if any, computer science work took place outside of CISE (maybe cryptography, networking, information/databases theory work, or computational mathematics)?

**[Gordon Bell]** Kent Curtis established Cryptography with NSA in 1987. Networking research got established too, and of course databases were part of CS research. So very little was done outside of CISE.

How did the work environment / morale / funding / power change with the switch from Division to Directorate?

**[Gordon Bell]** My biased perspective was that Morale improved and we had quite a lot of spirit by being together. I tried to break down parochialism that said money is where the power is.

Did the reorganization add topstige of Computer Science?

**[Gordon Bell]** I think so. The CS community complained vociferously that the supercomputing program shouldn't be part of CISE. I tried to get researchers to focus on parallelism and supercomputing as a source of ideas to enrich CS. This pretty much failed as they always wanted me to close the centers and give them more money for the same old stuff.

I believe getting a single advisory committee to help guide the direction of CISE helped the community.

What were the key topic areas CISE investigated/funded under your direction?

**[Gordon Bell]**

You should note this in the fax I sent. Parallelism was the big theme together with working on scientific computing... and networking!

How much of the work in the early years represented new areas of inquiry in contrast to continuation of research from earlier computing programs at NSF.

**[Gordon Bell]** LOTS: parallelism, Supercomputing, and NETWORKING, we talked about visualization but not that much happened then. We also bought some milling machines to get CS into robotics.

What were some of CISE's major successes?

**[Gordon Bell]**

The biggest success, by far, was part of the FCCSET program that was composed of all the research directors in the various agencies including DARPA that proposed the HPCCI that is still going strong. DOE was supposed to deal with supers. Darpa infrastructure and their program in high performance computing. I headed the interagency group that proposed the National Research and Education Network AKA NII AKA Internet in response to a bill from Senator Gore. We had a meeting in San Diego in the winter of 1987 where the proposal for a single research network came into being that NSF would run the attached slides show the plan I put forward to deploy bandwidth. Ultimately, I got Bob Kahn of CNRI to take on this program as an independent corporation to get around the bureaucracy and to help with operational issues. The first contract to operate the net went to Michigan and IBM.

Serendipity saved us because although we knew that bandwidth was key, we didn't see the kind of use that high bandwidth would enable. Gopher gave people a glimpse of a network and Behnres-Lee built www using http. Then the folks at U of IL made it ubiquitous. None of us saw that the real benefit of bandwidth would be to reduce latency so that anyone could interact with any machine anywhere provided you only wanted a few KB and could wait a few seconds.

(You can get my version of how this all came into being by looking at the slide show about Internet on my home page that is attached. Especially look at the jump in bandwidth in the attached slides that occurred

when we started the NSFnet!

How would you measure success?

**[Gordon Bell]** To really measure success would entail putting up some criteria. For Example, I considered running a piece of ARPA because it was better suited to my interests as an engineer to produce things. However, for NSF the measures are usually in terms of input... how much did we spend?

In the usual Washington bureaucratic fashion where input \$s are the measure, we succeeded. We got computing into a higher visibility position and got it on a more rapid budget track than it otherwise would have had if it had to compete for portions of MPS or as a part of engineering. Networking wouldn't have started and I can't imagine the separate government agencies getting together to do anything.

If you look at the increase in production of PhD's it improved enormously. This is NSF's usual measure... the actual breakthrough research that was funded that would otherwise not have been funded is unclear.

I must point out that NREN aka the NII is really the only interagency program that I know of. Agencies like doe, NASA, DOD/arpa are walled groups who have to maintain their constituency so doing anything collaborative is unique.

We initiated two distributed research centers programs during that time that I think have been interesting: graphics at brown, unc, cornell, and ??; and the supers center at cal tech, rice, etc. for software.

From my point of view, I would want to measure something great that we were able to get started in computing. Unfortunately, I'm not close enough to name names, but that's how I like to look at the world. Are we paying back the research investment? I would look at direct company and product technology transfer and I don't happen to know what these are.

I think we failed at being able to get a more meaningful supers program going with other directorates. Directorates, reflecting their compartmentalized university counter-parts are insular and don't want anything that crosses these boundaries. The only thing that makes the system work to do anything extraordinary are the occasional undergraduates and a few graduates that get something going among disciplines.

I see a failure as I just looked at a 1995 Computing Surveys on the CS computing agenda. It appears that it could have been written in 1985! CS has missed most of what has to do with graphics, and multimedia AKA video, etc. These are systems issues and it is hard to do these at NSF and universities because of grant size and personal risk to researchers.

I would also look at where the organization is. Did it's programs and direction stay in tact? (Good news and BAD news!)

Personally, I look at CISE in terms of whether I wasted my time, in terms of: did it get started right and have the right agenda, produce the usual PhDs and papers/\$, did we get at least a few great results from the researchers, and did it produce at least one really great result?

Would such measures be useful in today's environment given the requirements

of the Government Performance and Review Act (GPRA)?

**[Gordon Bell]** I'm really not sure what this is. I was asked to review something like this last year and it looked pretty useless, mixed up, and Washingtonized.



*Regards,*

**g**

Gordon Bell; 450 Old Oak Court; Los Altos, CA 94022; phone/fax 415 949 2735

<http://www.research.microsoft.com/barc/gbell>

**Possible BELL Questions: Internet History Archive Oral Interviews:**

2008-09-30

**Gordon Bell Questions**

1. Tell us about your childhood & childhood interests.
2. Tell us about your education & professional career.
3. What individuals were most influential in your career?
4. You are well known in the area of computing but tell us about your first experience with computer networks & networking?

How did your experience at DEC influence your view of networking?

5. You were AD at the NSF during 1986-87. Tell us about what happened during that period, particularly as it relates to networking.

Possible follow-up questions

- Relationship between supercomputing and networking
- Why TCP/IP protocols/NSFNET architecture
- 1987 backbone solicitation
- What factors in selecting Merit response
- Interesting anecdotes. (e.g. ATT pressure)
- NREN
- Political climate in DC

6. What were the motivations behind creating the NSFNET?

7. **What were the success factors in the creation of NSFNET?**

Who were your allies?

What were the barriers & impediments?

How did you over-come them?

8. Who do you consider the key people in the creation & operation of NSFNET & what did they contribute?

What happened to your collaborators?

9. What was your motivation?

What was the problem you were trying to fix?

10. What or who were the sources of technical ideas & innovations?

11. In the mid 1980s the network could have evolved in many directions. Why did it evolve the way it did?

12. If you were doing it over again, what would you do differently?

### **Possible BELL Questions: Internet History Archive Oral Interviews:**

13. Knowing what we know now, should the architecture or underlying technology of the Internet be implemented differently. Possible follow-up questions re security.

14. What did you then & what do you now view as government's role in cyber-infrastructure?

For science & technology?

15. Could the Internet be created today?

16. What do you see as the future of computing and networking?

How should it be funded?

1)

I'm a fan of crediting everything to my parents and environment that I grew up in. I was born in Kirksville, Missouri, on August 19th, 1934. Kirksville was a college town and also a farming community, about 10,000 population, and had the Northeast Missouri State Teacher's College, which has now morphed itself into the Truman University. The town now has about 20,000 people. My father had an electrical contracting business and appliance store and did small appliance repair, and I grew up working at "the shop". My uncle was with him and had the refrigeration part of the business. And so I spent all my formative years at the shop and in that environment.

**Bell:** I have a sister who is six years younger -- she's a school teacher and was a primary school principal at one point. My mother taught fourth grade and so was a teacher as well. So I grew up in that kind of environment. Mother didn't teach while I was growing up, though. We lived kind of on the outskirts of the town. We had a couple of acres and had a great garden, so I still remember what real tomatoes taste like and all those other things you get when you have your own garden. You can't get good tomatoes in California until mid August when the dry farmed ones come in.

**Hendrie:** What are the earliest memories you have of thinking about what you might want to do when you grew up?

**Bell:** That has been a really hard one for me to think of. I turned out to be one of the best dishwasher repair people in the area because it had cams and cycles and mechanical stuff then. I did house and building wiring, and at the time there was the REA or Rural Electrification Association or Administration. This was a federal program to wire all the farm houses in the country. So I went out and did a lot of wiring. I also installed industrial equipment and worked on all of that kind of stuff, e.g. appliances of all types, houses, buildings, and small industries like creameries, bottling plants, granaries.

In retrospect, I recall starting to earn a weekly salary of \$6 at some early age, perhaps when I was 12. This meant having a savings account and being comparatively rich. In those days of not having strict child labor laws, it was easy to work. My father felt that I was doing a man's work as an electrician at this age and should be paid accordingly. I do recall rewinding the stator of a motor when I was in the 5<sup>th</sup> grade. Today, you throw motors with burned out windings away.

**Bell:** In fact, I like to say I retired as a journeyman electrician when I went to MIT in 1952.

**Hendrie:** Well, now, tell me about the things you did, what you studied. You obviously learned a lot from just working with your father. What did you study in high school? Do you remember what courses you liked or...

### **Possible BELL Questions: Internet History Archive Oral Interviews:**

**Bell:** Yeah. What was really important was having a really wonderful science teacher and a wonderful math teacher. I still remember them very fondly. At that point in time and in Kirksville Missouri, you didn't take calculus or it wasn't offered, but I took chemistry and physics and then geometry and, maybe, solid geometry and foundational stuff. Those were really critical to enable me to go to MIT. At some point, I don't know when, maybe when I was 12 or so, I thought I wanted to be an engineer. I had no idea what an engineer was. I had books that I sort of read -- books of knowledge and an earlier kind of The Way Things Work -- so I gleaned that somewhere, somebody figured out how to make these things work and that was the interesting thing, not repair. Repairing them was okay but, in fact, designing them or inventing them seemed like a lot more fun. So that was basically the course that I set fairly early, with no one telling me I should or should not be doing this.

**Hendrie:** Or having some relative who was an engineer.

**Bell:** Yeah, I didn't have any.

2.

Discuss your personal history prior to joining DEC. Grew up in a college-farming town, Kirksville MO, worked in my father's electrical contracting and appliance biz; went to MIT got a master's as co-op student in 57, followed by a Fulbright to Oz. Came back worked for a year at MIT Speech Research Lab using the TX-0 that Ken, Andy, Ben Gurley, and Wes Clark had designed at Lincoln Lab. I was designing a tape unit controller for it, needed modules, and went to DEC to buy them.

- a) When did you join DEC? Got an offer May 24, 1960, from Harlan Anderson for \$9600 x 639.24% =61K. I joined in August 1960.
- b) How were you recruited? I was a customer, met everyone... Ben, Ken, Andy and they liked me and made me an offer. I met Dick Best, the Chief Engineer who was a great circuit engineer and really was responsible for getting us working components. Ben was a giant -- he was interesting and worked at all levels from devices to Operating Systems. Designed the analog displays including a high res scope.

In retrospect since these guys were great engineers, so how could I say no?

- c) Why did you join the company? It was what I thought engineering should be vs being part of the very large sea of desks environments I had seen as a coop student. I was the second engineer. It looked like I could really help and be the engineer I had dreamed of.

Describe your first interactions with Ken Olsen. Perhaps when I went to DEC to buy modules. I met Ken on maybe the 2<sup>nd</sup> or 3<sup>rd</sup> time. My interaction was with Ben Gurley and then Harlan Anderson who wrote the offer letter.

- d) What were your impressions? DEC had an environment that matched what I thought engineering should be: lots of responsibility to just design and build stuff. Interaction with customers to help get it right, etc. Smart people to interact with.
- e) Did you share the same vision about computing? I don't believe Ken and I have any shared vision of computing. Ken's view was as interesting artifact, like a car, to engineer. Ken may or may not have ever cared much about the actual use and programming.

My own view is more like children that you are helping grow. You have aspirations for them and work with each new generation to do something they have never done before. You feel they are the most important things in the world and the challenge is to help them reach their potential.

- f) Was there already a strong DEC culture in place? Yes. Much was: bi-weekly progress reports, open communication, The Engineering Committee and Engineering Notes that resulted in specifications, etc. It was non-hierarchical and politics free... you could talk to anyone, express opinions, consensus was good versus having to lobby a hierarchy to run it up the chain with formal decision making steps. This difference only becomes clear if you live in other environments. I like being first so having unique computers for unique apps was appealing. Given my speech experience, real time was important.

### Possible BELL Questions: Internet History Archive Oral Interviews:

3

**CGB 6/29/08:** *I have been incredibly talented, also called lucky, in finding people to work with or who have adopted me. Career claim: lucky & befriended, starting with my parents.*

- 1934 Parents providing a great home; allowing me to be the oldest person at MSFT;
- High school math Mrs. Strong and science teacher Bill Heinberg;
- 1952 MIT several extraordinary teachers & a few friends;
- 1958 Fulbright Scholar, Australia. Gordon Brown, an Aussie, who headed MIT EE Dept and meeting Gwen Druyor;
- 1959 MIT Speech Lab Ken Stevens, my thesis advisor and then I worked for him when coming back from Oz before going to DEC.
- 1960-83 DEC (#80) Ken Olsen, Ben Gurley, Harlan Anderson, Dick Best, Alan Kotok, Mary Jane Forbes, Dave Cutler & many others incl. Bill G;
- 1966-72 CMU Ivan Sutherland introduction, Al Newell, Al Perlis, Rod Williams, Herb Simon, Dan Siewiorek, Sam Fuller, Raj Reddy, ...'
- 1976-Computer Museum, Computer history Museum, Gwen Bell;
- 1983-2008 100+ startup entrepreneurs e.g. MIPS, SGI, NetApp, Cirrus Logic, John Sanguenetti, Steve Blank; Heidi Mason and I collaborated to create a methodology for evaluating startups.
- 1986 National Science Foundation (CISE) Eric Bloch; Chuck Brownstein was my Associate that knew you to operate the government.
- 1986-1989 Ardent Graphics Supercomputer, 35 great engineers;
- 1991 Microsoft Research—Nathan Myhrvold, Rick Rashid...;
- 1995 Microsoft Research—Jim Gray & Jim Gemmell & Vicki Rozycki; and Sheridan Forbes (significant other).

4.

You are well known in the area of computing but tell us about your first experience with computer networks & networking?

I was the project engineer for a message switching system and became enamored with telegraphy as the way to transmit information. In 1/61-11/61 I was the Project leader holding the systems responsibility for the DEC PDP 1 computer system marketed by ITT as the ADX-7300. Specific systems designs included: communications interface equipment for telegraph lines; a multiple channel interrupt for processing a large number (256) of lines; and duplex computer inter - connection control equipment.

I invented the UART Universal Asynchronous Receiver Transmitter. This allowed information to be sent in Teletype format first it was the 5 bit baudot code, then it evolved to 8 bit ASCII in mid 60s.

How did you experience at DEC influence you view of networking?

DEC had a very strong architecture that was based on DECnet, in a sense, a competitor to TCP/IP that we used to make switched peer2peer nets, unlike the IBM hierarchical nets using SNA.

Formatted: Underline

The DEC VAX Strategy architecture c1978 was predicated on networking! The VAX Architecture c1978 was a fully networked architecture. It was totally predicated on networking as every computer in a hierarchy going from mainframes to PCs was part of the network.

Formatted: Underline

And the culmination of that was I went to the board, I think in November or December of 78 or early 79, and I described what we were going to do and that became known as the VAX strategy. I used a diagram in the shape of an 'E'. At the top part of the 'E' was a cluster with a cluster interconnect and then all of these levels and things were connected together by an 'I' which was a network interconnect. And fortunately in the '78 time frame we were able to

### Possible BELL Questions: Internet History Archive Oral Interviews:

make the connection with Xerox, and Bob Metcalfe for the Ethernet.

**Hendrie:** Tell us a little bit about that story. What happened there?

**Bell:** Well, first off, the strategy was in a sense predicated on an 'I for Interconnect,' of having some way of doing that, and we had two or three different ideas about connecting everything using a local area network.

**Hendrie:** You had some idea of a way to do it but you hadn't thought of . . .

**Bell:** And at that time our friend Apollo wasn't started so we had two or three different network interconnect structures there that we were looking at. We had some rings and we had other stuff. We didn't know a lot about the Ethernet at that point although that had come out I guess in '76 or so and that's what PARC was using to connect their workstations. What happened was Bob Metcalfe joined us as a consultant. He had left PARC and said, "Gee, I'm here, what I want to do is propose that there be a standard here and I want to try to broker a deal here with Intel." Or I mean with Xerox. And so that was kind of the beginning of the DEC, Intel, Xerox cabal that in fact ended up proposing Ethernet. And I can probably find it, but I don't recall exactly the first meeting of that but it was with Phil Kaufman, who was later the president of . . .

**So basically DEC had three major components of today's networks and architectures:**

**a. DECnet, a forerunner to TCP/IP**

**b. we were the main driver for Ethernet in order to interconnect all computers, including using the computers themselves as switches. This is up there in importance as a component. Without it, the LANs wouldn't have developed.**

**c. a Cluster architecture for scalability so you could make big machines from more optimum or small machines**

Formatted: Font: Bold, Underline

Formatted: Font: Bold, Underline

Formatted: Font: Bold, Underline

**The DEC experience also taught me that networking was an entity into itself, not just connections that were bolted on to other systems! See later comments.**

Formatted: Font: Bold

5.

You were AD at the NSF during 1986-87. Tell us about what happened during that period, particularly as it relates to networking.

The first major decision I made was to create the networking division and to extricate it from the five Supercomputer centers groups who had the responsibility for their networks. Key to this was the hires. Dennis Jennings I believe led this and I was lucky to get Steve Wolff from Aberdeen who took it over. This was an essential part of networking because it gave seed money to build all the regional nets.

The second activity was the cross-agency group for networking that I chaired. Just as all the centers had their own networks, so too, each agency had one or two networks—NASA, DOE, DOD, an NSF. I was particularly sympathetic to ARPANet and getting them out of the operational aspects because it was eating their research budget. I had grown up both supplying computers from DEC and also being funded at CMU as a researcher. Also, I knew the ARPA folks, especially Steve Squires. The result of this was the NSF response to the Gore Bill. I believe that we had actually helped draft the bill that said put NSF in charge of proposing NREN.

We actually had to delay the submission of the proposal (3 volumes, see Volume 1) because there was an a government overlay committee that reported to Paul Hurray, who reported to Bill Graham, Pres. Reagan's Science advisor. The committee was called FCCSET and assigned

**Possible BELL Questions: Internet History Archive Oral Interviews:**

responsibilities like: NSF-networking, DOE-HPC and its apps, DARPA-infrastructure with people and also architectures for HPC. DARPA had their SCI program that was responding to the Japanese 5<sup>th</sup> Generation Challenge. In 1987 summer we had our report ready to go to congress proposing the net and my fellow committee members wouldn't let me submit it because it made their own efforts look amateurish. Eventually 19x?? Steve Squires was able to get our a nice comprehensive report that included all of the other parts, but by then, NSF was on its way to getting NREN going. The first bluebook of HPC came out in ???

Then the third part was actually getting Bob Kahn to manage NREN outside of the NSF bureaucracy.

Formatted: Superscript

**From:** Gordon Bell  
**To:** [REDACTED]  
**Subject:** FW: Government Oral History Invitation 10/2 at 9 AM (two FCCSET Papers)  
**Date:** Wednesday, October 1, 2008 6:41:00 PM  
**Attachments:** [080930-Bell Questions.doc](#)

---

The following comes from the word file of questions attached.

#### Gordon Bell Questions

1. Tell us about your childhood & childhood interests.
2. Tell us about your education & professional career.
3. What individuals were most influential in your career?
4. You are well known in the area of computing but tell us about your first experience with computer networks & networking?

How did you experience at DEC influence you view of networking?

5. You were AD at the NSF during 1986-87. Tell us about what happened during that period, particularly as it relates to networking.

#### Possible follow-up questions

- Relationship between supercomputing and networking
- Why TCP/IP protocols/NSFNET architecture
- 1987 backbone solicitation
- What factors in selecting Merit response
- Interesting anecdotes. (e.g. ATT pressure)
- NREN
- Political climate in DC

6. What were the motivations behind creating the NSFNET?
7. What were the success factors in the creation of NSFNET?

Who were you allies?

What were the barriers & impediments?

How did you over-come them?

8. Who do you consider the key people in the creation & operation of NSFNET & what did they contribute?

What happened to your collaborators?

9. What was your motivation?

What was the problem you were trying to fix?

10. What or who were the sources of technical ideas & innovations?
11. In the mid 1980s the network could have evolved in many directions. Why did it evolve the way it did?



12. If you were doing it over again, what would you do differently?
13. Knowing what we know now, should the architecture or underlying technology of the Internet be implemented differently. Possible follow-up questions re security.
14. What did you then & what do you now view as government's role in cyber-infrastructure?

For science & technology?

15. Could the Internet be created today?
16. What do you see as the future of computing and networking?

How should it be funded?

1)

I'm a fan of crediting everything to my parents and environment that I grew up in. I was born in Kirksville, Missouri, on August 19th, 1934. Kirksville was a college town and also a farming community, about 10,000 population, and had the Northeast Missouri State Teacher's College, which has now morphed itself into the Truman University. The town now has about 20,000 people. My father had an electrical contracting business and appliance store and did small appliance repair, and I grew up working at "the shop". My uncle was with him and had the refrigeration part of the business. And so I spent all my formative years at the shop and in that environment.

Bell: I have a sister who is six years younger -- she's a school teacher and was a primary school principal at one point. My mother taught fourth grade and so was a teacher as well. So I grew up in that kind of environment. Mother didn't teach while I was growing up, though. We lived kind of on the outskirts of the town. We had a couple of acres and had a great garden, so I still remember what real tomatoes taste like and all those other things you get when you have your own garden. You can't get good tomatoes in California until mid August when the dry farmed ones come in.

Hendrie: What are the earliest memories you have of thinking about what you might want to do when you grew up?

Bell: That has been a really hard one for me to think of. I turned out to be one of the best dishwasher repair people in the area because it had cams and cycles and mechanical stuff then. I did house and building wiring, and at the time there was the REA or Rural Electrification Association or Administration. This was a federal program to wire all the farm houses in the country. So I went out and did a lot of wiring. I also installed industrial equipment and worked on all of that kind of stuff, e.g. appliances of all types, houses, buildings, and small industries like creameries, bottling plants, granaries.

In retrospect, I recall starting to earn a weekly salary of \$6 at some early age, perhaps when I was 12. This meant having a savings account and being comparatively rich. In those days of not having strict child labor laws, it was easy to work. My father felt that I was doing a man's work as an electrician at this age and should be paid accordingly. I do recall rewinding the stator of a motor when I was in the 5th grade. Today, you throw motors with burned out windings away.

Bell: In fact, I like to say I retired as a journeyman electrician when I went to MIT in 1952.

Hendrie: Well, now, tell me about the things you did, what you studied. You obviously learned a lot from just working with your father. What did you study in high school? Do you remember what courses you liked or...

Bell: Yeah. What was really important was having a really wonderful science teacher and a wonderful math teacher. I still remember them very fondly. At that point in time and in Kirksville Missouri, you didn't take calculus or it wasn't offered, but I took chemistry and physics and then geometry and, maybe, solid geometry and foundational stuff. Those were really critical to enable me to go to MIT. At some point, I don't know when, maybe when I was 12 or so, I thought I wanted to be an engineer. I had no idea what an engineer was. I had books that I sort of read -- books of knowledge and an earlier kind of The Way Things Work -- so I gleaned that somewhere, somebody figured out how to make these things work and that was the interesting thing, not repair. Repairing them was okay but, in fact, designing them or inventing them seemed like a lot more fun. So that was basically the course that I set fairly early, with no one telling me I should or should not be doing this.

Hendrie: Or having some relative who was an engineer.

Bell: Yeah, I didn't have any.

2.

Discuss your personal history prior to joining DEC. Grew up in a college-farming town, Kirksville MO, worked in

my father's electrical contracting and appliance biz; went to MIT got a master's as co-op student in 57, followed by a Fulbright to Oz. Came back worked for a year at MIT Speech Research Lab using the TX-0 that Ken, Andy, Ben Gurley, and Wes Clark had designed at Lincoln Lab. I was designing a tape unit controller for it, needed modules, and went to DEC to buy them.

- a) When did you join DEC? Got an offer May 24, 1960, from Harlan Anderson for \$9600 x 639.24% =61K. I joined in August 1960.
- b) How were you recruited? I was a customer, met everyone... Ben, Ken, Andy and they liked me and made me an offer. I met Dick Best, the Chief Engineer who was a great circuit engineer and really was responsible for getting us working components. Ben was a giant -- he was interesting and worked at all levels from devices to Operating Systems. Designed the analog displays including a high res scope.

In retrospect since these guys were great engineers, so how could I say no?

- c) Why did you join the company? It was what I thought engineering should be vs being part of the very large sea of desks environments I had seen as a coop student. I was the second engineer. It looked like I could really help and be the engineer I had dreamed of.

Describe your first interactions with Ken Olsen. Perhaps when I went to DEC to buy modules. I met Ken on maybe the 2nd or 3rd time. My interaction was with Ben Gurley and then Harlan Anderson who wrote the offer letter.

- d) What were your impressions? DEC had an environment that matched what I thought engineering should be: lots of responsibility to just design and build stuff. Interaction with customers to help get it right, etc. Smart people to interact with.

- e) Did you share the same vision about computing? I don't believe Ken and I have any shared vision of computing. Ken's view was as interesting artifact, like a car, to engineer. Ken may or may not have ever cared much about the actual use and programming.

My own view is more like children that you are helping grow. You have aspirations for them and work with each new generation to do something they have never done before. You feel they are the most important things in the world and the challenge is to help them reach their potential.

- f) Was there already a strong DEC culture in place? Yes. Much was: bi-weekly progress reports, open communication, The Engineering Committee and Engineering Notes that resulted in specifications, etc. It was non-hierarchical and politics free... you could talk to anyone, express opinions, consensus was good versus having to lobby a hierarchy to run it up the chain with formal decision making steps. This difference only becomes clear if you live in other environments. I like being first so having unique computers for unique apps was appealing. Given my speech experience, real time was important.

3

CGB 6/29/08: I have been incredibly talented, also called lucky, in finding people to work with or who have adopted me. Career claim: lucky & befriended, starting with my parents.

- 1934 Parents providing a great home; allowing me to be the oldest person at MSFT;
- High school math Mrs. Strong and science teacher Bill Heinberg;
- 1952 MIT several extraordinary teachers & a few friends;
- 1958 Fulbright Scholar, Australia. Gordon Brown, an Aussie, who headed MIT EE Dept and meeting Gwen Druyor;
- 1959 MIT Speech Lab Ken Stevens, my thesis advisor and then I worked for him when coming back from Oz before going to DEC.
- 1960-83 DEC (#80) Ken Olsen, Ben Gurley, Harlan Anderson, Dick Best, Alan Kotok, Mary Jane Forbes, Dave Cutler & many others incl. Bill G;
- 1966-72 CMU Ivan Sutherland introduction, Al Newell, Al Perlis, Rod Williams, Herb Simon, Dan Siewiorek, Sam Fuller, Raj Reddy, ...'
- 1976-Computer Museum, Computer history Museum, Gwen Bell;
- 1983-2008 100+ startup entrepreneurs e.g. MIPS, SGI, NetApp, Cirrus Logic, John Sanguenetti, Steve Blank; Heidi Mason and I collaborated to create a methodology for evaluating startups.
- 1986 National Science Foundation (CISE) Eric Bloch; Chuck Brownstein was my Associate that knew you to operate the government.
- 1986-1989 Ardent Graphics Supercomputer, 35 great engineers;
- 1991 Microsoft Research—Nathan Myhrvold, Rick Rashid...;
- 1995 Microsoft Research—Jim Gray & Jim Gemmell & Vicki Rozycki; and Sheridan Forbes (significant other.

4.

You are well known in the area of computing but tell us about your first experience with computer networks & networking?

I was the project engineer for a message switching system and became enamored with telegraphy as the way to transmit information. In 1/61-11/61 I was the Project leader holding the systems responsibility for the DEC PDP 1 computer system marketed by ITT as the ADX-7300. Specific systems designs included: communications interface equipment for telegraph lines; a multiple channel interrupt for processing a large number (256) of lines; and duplex computer inter - connection control equipment.

I invented the UART Universal Asynchronous Receiver Transmitter. This allowed information to be sent in Teletype format first it was the 5 bit baudot code, then it evolved to 8 bit ASCII in mid 60s.

How did you experience at DEC influence you view of networking?

DEC had a very strong architecture that was based on DECnet, in a sense, a competitor to TCP/IP that we used to make switched peer2peer nets, unlike the IBM hierarchical nets using SNA.

The DEC VAX Strategy architecture c1978 was predicated on networking! The VAX Architecture c1978 was a fully networked architecture. It was totally predicated on networking as every computer in a hierarchy going from mainframes to PCs was part of the network.

And the culmination of that was I went to the board, I think in November or December of 78 or early 79, and I described what we were going to do and that became known as the VAX strategy. I used a diagram in the shape of an 'E'. At the top part of the 'E' was a cluster with a cluster interconnect and then all of these levels and things were connected together by an 'I' which was a network interconnect. And fortunately in the '78 time frame we were able to make the connection with Xerox, and Bob Metcalfe for the Ethernet.

Hendrie: Tell us a little bit about that story. What happened there?

Bell: Well, first off, the strategy was in a sense predicated on an 'I for Interconnect,' of having some way of doing that, and we had two or three different ideas about connecting everything using a local area network.

Hendrie: You had some idea of a way to do it but you hadn't thought of . . .

Bell: And at that time our friend Apollo wasn't started so we had two or three different network interconnect structures there that we were looking at. We had some rings and we had other stuff. We didn't know a lot about the Ethernet at that point although that had come out I guess in '76 or so and that's what PARC was using to connect their workstations. What happened was Bob Metcalfe joined us as a consultant. He had left PARC and said, "Gee, I'm here, what I want to do is propose that there be a standard here and I want to try to broker a deal here with Intel." Or I mean with Xerox. And so that was kind of the beginning of the DEC, Intel, Xerox cabal that in fact ended up proposing Ethernet. And I can probably find it, but I don't recall exactly the first meeting of that but it was with Phil Kaufman, who was later the president of . . .

So basically DEC had three major components of today's networks and architectures:

a. DECnet, a forerunner to TCP/IP

b. we were the main driver for Ethernet in order to interconnect all computers, including using the computers themselves as switches. This is up there in importance as a component. Without it, the LANs wouldn't have developed.

c. a Cluster architecture for scalability so you could make big machines from more optimum or small machines  
The DEC experience also taught me that networking was an entity into itself, not just connections that were bolted on to other systems! See later comments.

5.

You were AD at the NSF during 1986-87. Tell us about what happened during that period, particularly as it relates to networking.

The first major decision I made was to create the networking division and to extricate it from the five Supercomputer centers groups who had the responsibility for their networks. Key to this was the hires. Dennis Jennings I believe led this and I was lucky to get Steve Wolff from Aberdeen who took it over. This was an essential part of networking because it gave seed money to build all the regional nets.

The second activity was the cross-agency group for networking that I chaired. Just as all the centers had their own networks, so too, each agency had one or two networks—NASA, DOE, DOD, an NSF. I was particularly sympathetic to ARPAnet and getting them out of the operational aspects because it was eating their research budget.

I had grown up both supplying computers from DEC and also being funded at CMU as a researcher. Also, I knew the ARPA folks, especially Steve Squires. The result of this was the NSF response to the Gore Bill. I believe that we had actually helped draft the bill that said put NSF in charge of proposing NREN.

We actually had to delay the submission of the proposal (3 volumes, see Volume 1) because there was an a government overlay committee that reported to Paul Hurray, who reported to Bill Graham, Pres. Reagan's Science advisor. The committee was called FCCSET and assigned responsibilities like: NSF-networking, DOE-HPC and its apps, DARPA-infrastructure with people and also architectures for HPC. DARPA had their SCI program that was responding to the Japanese 5th Generation Challenge. In 1987 summer we had our report ready to go to congress proposing the net and my fellow committee members wouldn't let me submit it because it made their own efforts look amateurish. Eventually 19x?? Steve Squires was able to get our a nice comprehensive report that included all of the other parts, but by then, NSF was on its way to getting NREN going. The first bluebook of HPC came out in ???

Then the third part was actually getting Bob Kahn to manage NREN outside of the NSF bureaucracy.

-----  
I have started answering a few of the questions.  
The networking Vol. 1 report laid out the direction.  
Also attached is the overall FCCSET Committee Report for the other parts of HPC.  
g

-----Original Message-----

From: [REDACTED] [mailto:[REDACTED]]  
Sent: Tuesday, September 30, 2008 1:24 PM  
To: Gordon Bell; [REDACTED]  
Subject: Re: Government Oral History Invitation 10/2 at 9 AM

Gordon,

I have attached a list of possible questions. What usually happens is that we start out following the script and then the conversation takes on a life of its own. I am not particularly interested in facts and dates. That can be dug out of the written record. I am more interested in what people were thinking and why they made the choices that they did. What were the factors in making those decisions? The stuff that doesn't show up in the final report.

See you Thursday at 9,

Doug

>-----Original Message-----

>From: Gordon Bell [mailto:[REDACTED]]  
>Sent: Tuesday, September 23, 2008 04:52 PM  
>To: [REDACTED]  
>Subject: RE: Government Oral History Invitation 10/2 at 9 AM

>

>The lab is at the Powell Street BART station on Market if you come by BART.

>g

>

>Gordon Bell

>Microsoft Research

>835 Market St. Suite 700, San Francisco, CA 94103

><http://www.research.microsoft.com/~gbell>

>Cell: [REDACTED]

>

>

>-----Original Message-----

>From: [REDACTED]

>[mailto:[REDACTED]]

>Sent: Tuesday, September 23, 2008 1:50 PM

>To: Gordon Bell; [REDACTED]

>Subject: Re: Government Oral History Invitation 10/2 at 9 AM

>

>Your lab is fine. Is it the Market Street location? I will put some  
>questions together and send them to you before I hit the road Friday.  
>(We are making a detour to Yosemite on the way.)

>

>Doug

>>-----Original Message-----

>>From: Gordon Bell [mailto:[REDACTED]]

>>Sent: Tuesday, September 23, 2008 04:16 PM

>>To: [REDACTED]

>>Subject: RE: Government Oral History Invitation 10/2 at 9 AM

>>

>>This is fine.

>>Please send me any questions you can beforehand.

>>Do you want to meet at our lab?

>>Regards,

>>g

>>

>>Gordon Bell

>>Microsoft Research

>>835 Market St. Suite 700, San Francisco, CA 94103

>><http://www.research.microsoft.com/~gbell>

>>Cell: [REDACTED]

>>

>>-----Original Message-----

>>From: [REDACTED]

>>[mailto:[REDACTED]]

>>Sent: Tuesday, September 23, 2008 10:43 AM

>>To: Gordon Bell; [REDACTED]

>>Subject: Re: Government Oral History Invitation

>>

>>Would 9 pm Thursday Oct 2 work?

>>

>>Doug

>>>-----Original Message-----

>>>From: Gordon Bell [mailto:[REDACTED]]

>>>Sent: Monday, September 22, 2008 05:41 PM

>>>To: [REDACTED]

>>>Subject: RE: Government Oral History Invitation

>>>

>>>This would work. I am available (Oct 1) Wed PM, Th-Fri. So pick a time that you would like. Our office is 835  
Market at the Powell Street BART if you come that way.

>>>Regards,

>>>g

>>>Gordon Bell

>>>Principal Researcher, Microsoft Research

>>>Office: 835 Market St. Suite 700; San Francisco 94103

>>>Home: [REDACTED], San Francisco, CA 94111

>>>Office:415 972 6542; cell [REDACTED] home [REDACTED]

>>><http://research.microsoft.com/~gbell>

>>>

>>>

>>>-----Original Message-----

>>>From: [REDACTED]

>>>[\[mailto:\[REDACTED\]\]](mailto:[REDACTED])

>>>Sent: Monday, September 22, 2008 12:32 PM

>>>To: Gordon Bell; [REDACTED]

>>>Subject: Re: Government Oral History Invitation

>>>

>>>Gordon,

>>>

>>>I will be in SF area (Burlingame) next week. Would you be available for a couple of hours for a recorded interview Tuesday-Friday. With the material from the NSFNET anniversary we now have over a hundred hours of material and a waiting to to hear if we were successful in getting an NSF grant to move the archive into production status. Charles Brownstein has passed on a number of "Gordon Bell" stories (not recorded) and I would like to get them from straight from you.

>>>

>>>Doug

>>>>-----Original Message-----

>>>>From: Gordon Bell [\[mailto:\[REDACTED\]\]](mailto:[REDACTED])

>>>>Sent: Monday, October 29, 2007 11:48 AM

>>>>To: [REDACTED]

>>>>Subject: RE: Government Oral History Invitation

>>>>

>>>>Great. Either work.

>>>>g

>>>>

>>>>-----Original Message-----

>>>>From: [REDACTED]

>>>>[\[mailto:\[REDACTED\]\]](mailto:[REDACTED])

>>>>Sent: Monday, October 29, 2007 8:18 AM

>>>>To: Gordon Bell; [REDACTED]

>>>>Subject: Re: Government Oral History Invitation

>>>>

>>>>I have ordered some high end equipment that claims to do enough digital signal processing to allow for a good recording over a phone line---but I haven't tried it out. I travel to SF frequently to visit grandkids (and parents too) so a face to face might be better. Let me talk to my wife and figure out when we will next be in SF.

>>>>

>>>>Doug

>>>>>-----Original Message-----

>>>>>From: Gordon Bell [\[mailto:\[REDACTED\]\]](mailto:[REDACTED])

>>>>>Sent: Monday, October 29, 2007 10:51 AM

>>>>>To: [REDACTED]

>>>>>Subject: RE: Government Oral History Invitation

>>>>>

>>>>>Happy to be interviewed here.

>>>>>Let me know when that would work and we can schedule.

>>>>>Could it be by phone?

>>>>>G

>>>>>

>>>>>Gordon Bell

>>>>><http://www.research.microsoft.com/~gbell>

>>>>>Principal Researcher, Microsoft Research

>>>>>Work: 455 Market Street, Suite 1690; San Francisco, CA 94105

>>>>>Homes: [REDACTED], San Francisco, CA 94111 Aurora

>>>>>Place, [REDACTED]; Sydney, NSW 2000

>>>>>Phones: [redacted] cell [redacted]  
>>>>> [redacted] home [redacted]  
>>>>>  
>>>>>  
>>>>>

>>>>>-----Original Message-----

>>>>>From: [redacted]  
>>>>>[mailto:[redacted]]  
>>>>>Sent: Monday, October 29, 2007 6:26 AM  
>>>>>To: Gordon Bell; [redacted]  
>>>>>Subject: Re: Government Oral History Invitation  
>>>>>

>>>>>Sorry to hear that Gordon and thanks for the attachment. Would you be available for a personal interview?  
If memory serves me right you are in the San Francisco area. Is that correct.

>>>>>

>>>>>Doug

>>>>>-----Original Message-----

>>>>>>From: Gordon Bell [mailto:[redacted]]  
>>>>>>Sent: Friday, October 26, 2007 08:01 PM  
>>>>>>To: [redacted]  
>>>>>>Subject: RE: Government Oral History Invitation  
>>>>>>

>>>>>>I'm sorry I can't attend.

>>>>>>

>>>>>>Below is from a Smithsonian Interview relative to my role while at NSF.

>>>>>>

>>>>>>Also attached is the FCCSET group report I chaired that proposed NREN!

>>>>>>

>>>>>>Ironically, I think the most important thing in the report is the GRAPH that established a timetable and plan.

>>>>>>

>>>>>>I believe it is about the only cross agency plan to come from Washington.

>>>>>>

>>>>>>

>>>>>>

>>>>>>Networking today starting with Internet 2 seems to me to be a complete disaster!

>>>>>>

>>>>>>G

>>>>>>

>>>>>>Gordon Bell

>>>>>><http://www.research.microsoft.com/~gbell>

>>>>>>Principal Researcher, Microsoft Research

>>>>>>455 Market Street, Suite 1690

>>>>>>San Francisco, CA 94105

>>>>>>Phone [redacted] cell [redacted]

>>>>>>

>>>>>>

>>>>>>

>>>>>>

>>>>>>NSFnet transition to NREN and the modern Internet The main thing I

>>>>>>did that I think was really important concerned the NSFNet and how

>>>>>>it became NREN or the Internet we have today. The net was

>>>>>>established as part of and reported to the person who ran the supercomputer centers division. I came in and  
said:

>>>>>>“Networking is going to report directly to me as a new division

>>>>>>and not to the supercomputer division. The network is independent and distinct from the supercomputer  
centers”

>>>>>>10 The original centers included: the University of Minnesota,

>>>>>>UC/San Diego, University of Illinois, the Pittsburgh Center, Princeton, and Cornell. In 2000 there were two at UC/San Diego and Illinois.

>>>>>>

>>>>>>This was based on my experience at DEC. The other part of the VAX  
>>>>>>Strategy was that we had built super networking technology called  
>>>>>>DECnet, by having a network group. It wasn't part of the computer  
>>>>>>guys who said: "We'll simply put UARTs in our computers and  
>>>>>>connect them to each other – we'll do the networking." Where is  
>>>>>>the network and why do we need a group to make links? Well the network is all of those lines and links, and  
>>>>>>it's especially all the code that makes the collection of computers work as one. So I did the same thing and said:  
>>>>>>"NSF needs a strong, independent networking group. We're going to  
>>>>>>build a network. We're starting all of that."

>>>>>>

>>>>>>And so I'd say I am most comfortable with my Washington experience leading networking.  
>>>>>>We said we were going to take a lead position, the Gore Bill came  
>>>>>>out in 1986, and NSF was given the charter to lead the group on  
>>>>>>networking across all the government agencies. And then again I  
>>>>>>would like to say that NREN (for National Research and Education  
>>>>>>Network) is the only thing I can cite as inter-agencies ever doing  
>>>>>>together and agreeing on. We got everybody together from all  
>>>>>>government agencies, industry, and academe and put a plan forward in February of 1987, that was a three-  
>>>>>>phase plan to provide bandwidth. And why this is really fresh is I gave a keynote talk at InternetWorld '95 in April.

>>>>>>

>>>>>>It's the role of serendipity. Most everyone think that the  
>>>>>>Internet just happened overnight. But it didn't. We had a three-day workshop of 500 people in San Diego  
>>>>>>talking about networking.

>>>>>>We had industry - what's bandwidth going to be like? All the  
>>>>>>government agencies - what are the needs? To the academics - what can you do?

>>>>>>

>>>>>>On the final morning, after listening to the previous two days,  
>>>>>>another "aha" occurred that was fundamentally the NREN plan. I drew it on a single overhead that everyone  
>>>>>>understood.

>>>>>>Figure 2. Plan for NREN created at the February 1987, San Diego NSF sponsored meeting.

>>>>>>I basically said: "Here's the plan. We're really have nothing now.  
>>>>>>Our networks are overloaded and really don't work very well. Phase  
>>>>>>Zero: We get ourselves together. We make the network solid. So  
>>>>>>without a system running no one is going to believe you about the  
>>>>>>future. Then we're going to go from 56 kilobit's today in the  
>>>>>>backbone to 1.5 megabit's in 1990 using T1 and then we go immediately to 45 Mbits. In 96-97 we'll start to  
>>>>>>field test the first gigabit nets. The later stage is research, the earlier network is strictly engineering."

>>>>>>

>>>>>>I called them Internet 1, 2, and 3 in a recent talk that I  
>>>>>>keynoted at InternetWorld 1995. One is ARPAnet, running 56-kilobit  
>>>>>>prototype for email. Two is what we've got today, which was really  
>>>>>>mail as a reliable delivery, the worldwide web and a prototype for  
>>>>>>three. And here's what three is: telephony, audio, video, and video conferencing. It can't be ubiquitous  
>>>>>>without fiber optic speeds, there's not enough capacity. And that's three to five years down the pike.  
>>>>>>Meanwhile we can have a lot of fun with what we're doing with Internet 2.

>>>>>>

>>>>>>Interestingly, the goal of ARPAnet was not mail, mail was not even  
>>>>>>conceived of. It was remote log in to other systems and sharing  
>>>>>>files. The plan didn't say anything about the application in our  
>>>>>>goals, we didn't say anything about worldwide web. We had no idea.  
>>>>>>It was proposed to be used for supercomputing. Well, all the  
>>>>>>networkers knew it wasn't supercomputers. There was no demand. We  
>>>>>>knew that supercomputers needed bandwidth, they needed to  
>>>>>>communicate, but when you really force people to use them they



>>>>>>would prefer their own machines. I talked to various folks at DOE  
>>>>>>about this dilemma. If you really want to get a lot of power  
>>>>>>together why don't you have Los Alamos run all your computers.  
>>>>>>You've got plenty of power, you have it together. The networking  
>>>>>>is just fine. In supercomputing there is no reason to have more  
>>>>>>than one computer in the center of the earth. In fact there is  
>>>>>>every reason not to except for the de-attachment you get. You get  
>>>>>>some attachment of these people coming together. Leading the NREN  
>>>>>>effort across all the agencies that created the  
>>>>>>  
>>>>>>network plan was the other thing I did at NSF I'm proud of.  
>>>>>>  
>>>>>>  
>>>>>>

>>>>>>-----Original Message-----

>>>>>>From: [REDACTED]  
>>>>>><[mailto:\[REDACTED\]](mailto:[REDACTED])>  
>>>>>>Sent: Friday, October 26, 2007 3:00 PM  
>>>>>>To: [REDACTED];  
>>>>>>[REDACTED];  
>>>>>>[REDACTED]; Gordon Bell;  
>>>>>>[REDACTED];  
>>>>>>[REDACTED];  
>>>>>>Cc: [REDACTED]  
>>>>>>Subject: Government Oral History Invitation  
>>>>>>  
>>>>>>  
>>>>>>  
>>>>>>

>>>>>>Colleagues,  
>>>>>>  
>>>>>>  
>>>>>>

>>>>>>As you may know, one of my pet projects over the past few years has been developing an archive of materials relating to the creation of the Internet. That archive now includes a number of oral history interviews as well as source documents. I am pleased that we have obtained funding to do eight professionally moderated and recorded group oral histories on Wednesday November 28th, the day before the NSFNET 20th Anniversary event. These group oral histories will augment the individual oral histories already collected and take advantage of the group dynamics that often trigger forgotten memories. A complete description of this "pre-Anniversary event" is attached.  
>>>>>>  
>>>>>>  
>>>>>>

>>>>>>Because of your collective leadership in nurturing the NSFNET and Internet, we would like to invite you to participate in the "Government" oral history group. The history of the ARPANET has been well documented and told. However, the transition from a specialized research network to a general purpose Internet is neither well documented nor well understood--yet I believe that it holds many important science policy lessons. The government played a crucial role in this transition. Your stories and memories are important if future scholars are to fully understand the formation of what is now regarded as one of the signature technological achievements of the 20th century.  
>>>>>>  
>>>>>>  
>>>>>>

>>>>>>The government session will be held from 1 pm to 4 pm at the conference hotel. I would appreciate a quick response regarding your participation. We are limiting the size of the groups and want to be sure that we have included as many key contributors as possible.  
>>>>>>  
>>>>>>

>>>>>

>>>>>Thanks,

>>>>>

>>>>>

>>>>>

>>>>>Doug Gale

>>>>>

>>>>>

>>>>>

>>>>>

>>>>>

>>>>>

>>>>

>>>>

>>>>

>>>

>>>

>>>

>>>

>>

>>

>>

>>

>

>

>

>

**From:** andreu@veabaro.info  
**To:** [Gordon Bell](#)  
**Subject:** RE: Comments on the NSF-led initiative for NREN and Internet  
**Date:** Sunday, November 7, 2004 4:38:21 PM

---

Thanks a lot. I already had integrated your graph in your interview.  
This comments are important and arrive at the perfect time.

andreu

---- Mensaje Original ----

De: [REDACTED]  
Para: andreu@veabaro.info  
Asunto: RE: Comments on the NSF-led initiative for NREN and Internet  
Fecha: Sun, 7 Nov 2004 10:10:32 -0800

>  
>Andreu,  
>Some more comments on the NSF NREN plan that I wrote in April 2000 to  
>an  
>interviewer.  
>  
>From a perspective of you as a visualization person, the graph of  
>bandwidth(t) plan for NREN aka Internet is probably the best way to  
>characterize what I think I did.  
>  
>Also see the attached interview re this and future apps re  
>Telepresence.  
>g  
>  
>> <<FCCSET 1987 Vol 1c.pdf>>  
>>  
>> I am proud of three things as the first, head of NSF's Computer and  
>> Information Science and Engineering Directorate. Erich Bloch hired  
>me  
>> (Erich, as you know is one of the greats in computing!) in 1986  
>> to start up the new directorate.  
>> \* Starting the directorate and setting the agenda that was focused  
>on  
>> parallelism. I took a modest amount of flak from the cs  
>> community including Ken Thompson and Don Knuth.  
>> \* Taking all the networking responsibility away from the  
>supercomputer  
>> centers, recruiting Steve Wolff and heading the cross-agency  
>> committee that  
>> produced that attached report that responded to the Gore Bill for  
>NREN. I  
>> was the NSF representative (heading the network part) to the  
>> FCCSET committee that included DOD (supers), DARPA (Saul Amarel for  
>> Infrastructures and basic CS).  
>> BTW: I outlined a much more ambitious structure that I couldn't  
>get  
>> the Science Advisor, Bill Graham to listen to. I have it  
>> if you want to see it. It is much like PITAC, but government  
>centered.

>>  
>> \* The attached report was a result of a very large 3 day meeting  
>of  
>> gov, industry, and academe in Feb. 1987 in San Diego. I sat in on  
>the  
>> various groups and on the morning of the last day, I drew the  
>diagram  
>on  
>> page 11 of the report that litterally pulled everything to gather  
>as a  
>> rallying plan. Ironically, it is the best schedule I've ever made  
>because  
>> it accurately went from '87-2000. The upshot of the report was  
>that  
>just  
>> as we were about to submit the report to respond to the Gore Bill,  
>my  
>> friends on FCCSET essentially said: "you can't do that, it makes us  
>look  
>> like idiots. Let us put this all together and we'll ge a shit load  
>of  
>> money." So that was the genesis of the first blue book that was  
>put  
>> together.  
>>  
>> I site the report and the ensuing work as the ONLY successful  
>cross-agency  
>> work. In part it was due to the simple one page plan for bandwidth.  
>>  
>> You might want to use what I see as a virtuous cycle of bandwidth:  
>> Increasing capacity > increases bandwidth > decreases latency >  
>> creates new apps > creates demand > increases capacity....  
>> This explains the transition from ftp to email to gopher and www to  
>audio  
>> and to video  
>> \* Interaction with supercomputing centers since this was part of  
>CISE.  
>> The irony here was that the first thing the directors told me was  
>that  
>VAX  
>> destroyed supercomputer centers and that I had a conflict of  
>interest  
>> because I liked small, personal machines. Then I took their  
>networks  
>away  
>> that they were all building as stars. Next I told them to use  
>> UNIX and they  
>> weren't going to help DOE maintain a timesharing system for their  
>Crays. I  
>> said they needed to try some of the new scalables that were  
>> coming from the  
>> SCI program. Also, I told Ken Wilson, Nobel Laureat that he had to  
>deliver  
>> training and cycles and was not permitted to do research for the  
>FPS T  
>> machine. Finally, I stopped building new centers and wouldn't  
>> pay ETA until

>> they delivered their computer... which was never. The result was to  
>start  
>> reducing the no. of centers from 6 to now 2.  
>>  
>

**From:** Leonard Kleinrock  
**To:** [Gordon Bell](#)  
**Subject:** Re: NSF Networking: Bell article Feb 88 on the NREN Plan that was started; 3 Vol NSF Response for the Gore Bil and Feb. 1987 Workshop; Finaly HPC report Nov. 87.;  
**Date:** Tuesday, December 2, 2014 8:35:23 PM

---

Right you are. But I don't consider Isaacson a good historian for many of the reasons you cite. I do agree that separation from events by many years (lifetimes) makes it very hard for historians to get it right.

So, given that the historians don't have the proper tools, what a great research and development project to support the creation of standard tools and sensible metrics and tests for accuracy and completeness (hard to do, but certainly worthwhile).

> On Dec 2, 2014, at 7:54 PM, Gordon Bell <[REDACTED]> wrote:

>

> The problem is that the documents have to be read and put in a clear time table.

> Historians don't start with a great database and timeline tool.

>

> The problem is history is basically a story with the rule that you can't manufacture facts ... regardless of their source...

> BUT you can leave out anything that screws up the story or effects the patina of the story's characters.

>

> An interesting example is the massive Isaacson history or collection of stories (conflicts, everything but sex) -- a compendium of heroes and controversial stories: Ada; Eckert- Mauchly - Atanasoff; Kilby-Noyce Moore; arguments among the networking and internetworking fathers--thankfully Tim just wrote his own code; Engelbart-Kay Parc and Metcalfe; two Steves; Gates-Allen-Kildall-IBM; Linus sans Thompson-Ritchie. Forget IBM, Cray; etc. because the book was too long.

> g

>

> -----Original Message-----

> From: Leonard Kleinrock [[mailto:\[REDACTED\]](#)]

> Sent: Tuesday, December 2, 2014 7:25 PM

> To: Gordon Bell

> Subject: Re: NSF Networking: Bell article Feb 88 on the NREN Plan that was started; 3 Vol NSF Response for the Gore Bil and Feb. 1987 Workshop; Finaly HPC report Nov. 87.;

>

> Yes and no. The good historians weigh all the evidence and facts. That's why it is so important to get these interviews archived.

>

>> On Dec 2, 2014, at 7:15 PM, Gordon Bell <[REDACTED]> wrote:

>>

>> History is created by the last guy standing.

>> g

>>

>> -----Original Message-----

>> From: Leonard Kleinrock [[mailto:\[REDACTED\]](#)]

>> Sent: Tuesday, December 2, 2014 7:13 PM

>> To: Gordon Bell

>> Subject: Re: NSF Networking: Bell article Feb 88 on the NREN Plan that was started; 3 Vol NSF Response for the Gore Bil and Feb. 1987 Workshop; Finaly HPC report Nov. 87.;

>>

>> Thanks for a great interview today, Gordon. And thanks for the attachments. It is quite a history and I hope the historians get it right.

>>

>> Regards to you and Sheridan,

>> Len

>>

>>> On Dec 2, 2014, at 4:15 PM, Gordon Bell <[REDACTED]> wrote:

>>>

>>> Len,

>>> Attached are several documents re. NSF and NREN. Note that the Michigan-IBM network was up and running by the end of 87 as I refer to it in my Spectrum paper, Feb. 88

>>>

>>> Note the dates...

>>> Cover for volume one and reference to the Gore bill. Three volumes attached. Vol II is summary of Workshop Feb 17-19, 1987.

>>> The plan for the Internet in terms of the bandwidth, I drew on the last morning, Feb 19 that was actually implemented is given below.

>>> I sketched it on an overhead which I do not have.

>>>

>>> BTW: Chuck Brownstein, my AD at NSF and then Kahn's Exec Officer would be worth interviewing! He stayed on at NSF.

>>> Bill Wulf took over for me at NSF.

>>>

>>> Finally, since the networking report was so clear, simple and strong, my colleagues were embarrassed and wouldn't let me submit it to Bill Graham the President's (Reagan) Science Advisor. His admin was Paul Huray.

>>>

>>> The FCCSET committee used the specificity of Network as the carrier for the HPC proposal that were published for a number of years (the first being attached).

>>> Squires and Bill Scherlis put the report together.

>>> The first HPC Proposal is attached. The Networking report was the only one with any goals and specific budget.

>>> g

>>>

>>> <image001.png>

>>>

>>> Here's the plan I drew in Feb. 87 after the meeting.

>>>

>>> <image002.png>

>>>

>>> <Bell\_Calls\_for\_US\_Research\_Network\_IEEE\_8802.pdf><FCCSET Report, Study of Critical Problems and Future Options, Reports fr....pdf><FCCSET Report, Study of Critical Problems and Future Options, Compendium....pdf><FCCSET Report, Study of Critical Problems and Future Options, Recommenda....pdf><FCCSET Research and Development Strategy for High Performance Computing ....pdf>

>>

>

**From:** Gordon Bell  
**To:** [Leonard Kleinrock](#)  
**Cc:** [Bradley Fidler](#); [Morten Bay](#)  
**Subject:** RE: Questions from Kleinrock email 2014-11-19  
**Date:** Wednesday, November 19, 2014 4:55:24 AM

---

The 2nd of December works for me.

Sent from my Windows Phone


---

**From:** [Leonard Kleinrock](#)  
**Sent:** 11/18/2014 11:19 PM  
**To:** [Gordon Bell](#)  
**Cc:** [Bradley Fidler](#); [Morten Bay](#)  
**Subject:** Re: Question

Hi again Gordon,

Do let me know if Dec 2 works for you.

Thanks,  
Len

> On Nov 17, 2014, at 5:00 PM, Leonard Kleinrock < wrote:

>

> Gordon,

>

> Thanks for the detailed reply. By way of background, the folks at today's DARPA are interested in learning what it was about the ARPA of the 1960's and 1970's that led to so many great accomplishments and successes.

>

> Specifically, they have contracted me to establish an interview series of a few key individuals who worked on the high-risk, high-payoff ARPANET project. You were part of it and are an excellent candidate and I hope you will consider being interviewed.

>

> This effort is not about assigning credit for personal accomplishments. Instead, it is about illuminating and articulating the early successful institutional culture and innovation strategy at DARPA, and thus provide insights that may be useful for the agency now and in the future.

>

> We would love to do a relatively short (say 1 to 1.5 hour) video interview with you to discuss the general issues of ARPA funding culture back then and also how ARPA helped or hindered the transition to a TCP/IP NSFNET as the internet backbone. We will address some other general issues as well.

>

> The ideal time for us would be on Dec 2, Tuesday, the day before our TTI/Vanguard Board meeting in SF. I suggest we use the suite I will get for my (and Stella's) stay at the Grand Hyatt San Francisco (at Union Square). We should start at 8 am since we are interviewing Norm Abramson at 11 am same location. I do hope that works for you.

>

> Do let me know.

>

> I am attaching some of the boilerplate docs that inform you about the University procedures, etc. We will bring some docs for you to sign when we see you.

>

> Best regards,

> Len



>  
> <Basic Legal Agreement.pdf><Use Policy.pdf><Introduction Letter 201406.pdf>

>  
>  
>> On Nov 17, 2014, at 11:03 AM, Gordon Bell <[REDACTED]> wrote:

>>  
>> I was funded at CMU 1966-72 and attended one of the meetings. I suggested (to Wes and ??) using a single byte seridal interface to connect IMPs to various mainframes and looked over that spec... that I assume was implemented. Also, they funded my architecture work that ultimately included book, several computer designs including C.ai that got built at Livermore as a 16 processor machine.

>>  
>> Yes, worked with ARPA/Darpa Squires and Bill Scherlis in making the transition that came from the NREN proposal that was driven from NSF.

>> There was an overall FCCSET Committee that reported to Paul Huray and Bill Graham (chief scientist I believe under Reagan), both appointees.

>>  
>> FCCSET committee was DARPA (Saul Amarel) representing Computer Science infrastructure funding, DOE (Jim Decker) responsible for HPC, me (networking).

>> We agreed that we were going to have an High Performance Computing initiative that covered everything.

>> The networking section was the only one that had any substance and Huray and the committee didn't want the network report (Vol 1) that we did to be published alone because it was the only strong one and the other guys hadn't finished their ... and they didn't have much to say other than fund CS research algorithms, and fund supercomputers and fund people (all). In the end both were published, but the networking carried the report, while HPC in the title carried the marketing pizzas.

>>  
>> The report plus the NSF network report that I sent around all 3 volumes, Vol 1 came in ahead of time to respond to the Gore Bill that NSF helped author.

>>  
>> Anyway, these two reports are pretty much the plan. Ironically, the networking is the only one that really got implemented because it was a plan.

>> Squires and Scherlis and Chuck Brownstein (NSF) drove to get the FCCSET HPC report done. Both came out concurrently.

>> g

>> -----Original Message-----

>> From: Leonard Kleinrock [[mailto:\[REDACTED\]](mailto:[REDACTED])]

>> Sent: Sunday, November 16, 2014 10:44 PM

>> To: Gordon Bell

>> Cc: Bradley Fidler; Morten Bay

>> Subject: Question

>> Hi Gordon,

>>  
>> I have a DARPA contract to study the ARPA funding culture that led to the great achievements of the 1960's and 1970's, etc. We have done some video interviews of some of the PI's, grad students, PM's, IPTO directors, DARPA directors.

>> We are using the ARPANET development as our prime example of a successfully funded ARPA project. I wonder if you had ARPA/DARPA funding during that period (or ever). And, importantly, I know you were involved with NSF taking over the Internet backbone in the form of NSFNET in the 1980's; my question is, were you involved with the interaction/negotiation/deliberation between NSF and DARPA for the NSFNET transition? I ask because if so, it may be appropriate for us to conduct a video interview with you while we are at the TTI/Vanguard meeting in SF early next month (probably early morning of Tuesday Dec 2 in my hotel room at the Grand Hyatt Union Square); this would make sense if you were involved with the DARPA/NSF deliberations.

>>  
>> Do let me know. In any case, looking forward to seeing you in SF for the TTI/Vanguard meeting.

>>  
>> Much thanks,

>> Len

>>

>> PS I am cc'ng two key folks working with me on this project.

>> <FCCSET Research and Development Strategy for High Performance Computing 871120 c.pdf><FCCSET Report, Study of Critical Problems and Future Options, Recommenda....pdf>

[REDACTED]

---

**From:** [REDACTED]  
**To:** 'William Aspray'  
**Cc:** 'Peter A. Freeman'; 'Rick Adrion'  
**Subject:** RE: CISE Oral History: 20 Questions  
**Attachments:** FCCSET Research and Development Strategy for High Performance Computing 871120 c.pdf

William,  
More.  
G

---

**From:** Gordon Bell  
**Sent:** Thursday, July 13, 2017 10:23 AM  
**To:** 'Gordon Bell' <[REDACTED]> 'William Aspray' <[REDACTED]>  
**Cc:** 'Peter A. Freeman' <[REDACTED]> 'Rick Adrion' <[REDACTED]>  
**Subject:** RE: CISE Oral History (1-11)

See comments below.

I sent docs on nren.

The other biggie was that the ASC i.e. supers program direction, funding:

use UNIX, stop von Neumann Center because CDC can't deliver, get the directorates to pay. Got Cray to support the centers, also IBM at Cornell

g

---

**From:** Gordon Bell  
**Sent:** Thursday, July 13, 2017 8:48 AM  
**To:** 'William Aspray' <[REDACTED]>  
**Cc:** Peter A. Freeman <[REDACTED]> Rick Adrion <[REDACTED]>  
**Subject:** RE: CISE Oral History

Fine.

g

---

**From:** William Aspray [[mailto:\[REDACTED\]](mailto:[REDACTED])]  
**Sent:** Thursday, July 13, 2017 8:44 AM  
**To:** Gordon Bell <[REDACTED]>  
**Cc:** Peter A. Freeman <[REDACTED]> Rick Adrion <[REDACTED]>  
**Subject:** CISE Oral History

Gordon,

I will plan to call you at 10:30 am PDT today at [REDACTED]

In case it is more convenient for you to start earlier, I am prepared and available. Just send me an email telling me when. If I don't hear from you, I will call as originally planned.

Below are the questions I plan to ask in case you want to think about them before the interview.

Thanks for agreeing to this interview.

Bill

\*\*\*

**Gordon Bell/CISE History – Interview Questions [likely to be followed up with other questions based upon comments made by the interviewee] [13 July 2017]**

- 1) Before joining CISE, what experience had you had with NSF more generally (grantee, reviewer, advisory board, etc.)? What about with other federal agencies (DARPA, ONR, etc.)?

I had an NSF Grant or two at CMU during 1966-72. Had served on various panels including the first one that reviewed centers proposals... vividly recall a Xerox researcher rejecting the Santa Clara U magnetics proposal by various former IBM disk folks because who was the university and that magnetic disks were dead with optical stores. Was on an industry panel that got a bill passed that would allow companies to gift stuff at a tax advantage that had a nice effect.

Also see attached letter.

- 2) Who recruited you to CISE? Why did you decide this was a good offer to accept? When did you actually arrive?

Erich Bloch. I knew Erich and NSF and thought it was an important thing to do.. also I wanted to work for Erich.

I felt strongly that computing should separate from being distributed appendages in other directorates and to have its own directorate!

I believe I arrived Jan 86... but need to check the date.

- 3) Was the fact that you had both high-level academic and high-level industrial experience an asset at NSF? Was the fact that you were distinguished in your technical career an asset in carrying out your NSF work?

Hopefully, my experience with large organizations, university faculty, and especially the computing industry re. what I felt was needed was useful

As you see in one of the interviews, I pressed the community for working on parallelism and got pushback from Knuth, Karp and Ken Thompson. Re dictating to community.

Failed to get computational science adopted by computer science... but did establish the name.

Kent Curtis had funded Seitz at Cal Tech. I got Darpa to pick this up and this was the path that got Intel et al into building MPPs, etc. for HPC

Also the connection with Squires et al at Darpa was important

- 3) Tell me about your relationship prior to and while at NSF with Eric Bloch?

I had massive respect for Erich and had worked with a committee he and Bob Noyce chaired to set up Semetech and SRC.

Did your shared engineering background with Bloch stand you in good stead at NSF, where science rather than engineering was pre-eminent?

Erich definitely establish engineering as an equal part Suh Nam, former MIT ME Dept head was a giant that Erich headed for the Engineering Directorate. He is president now of KEIST. Staff meetings were fun.

- 4) Although it predated you, what can you tell me about the reasons behind and the process involved in forming CISE?

The obvious branches were brought over. I had known Kent Curtis, Erich made decision that Bernie Chern would come over from Engineering, John Connolly from ASC, and ?? came over from robotics. Erich made the CISE name. I preferred "Computing" but we didn't spend a lot of time on the name.

No doubt, Chuck Brownsteing as my exec/adm assistant was perhaps the most important addition since he knew the various people and politics.

The big change I made fairly early on was to extract networking from the centers program. Am not sure about the details, but it was one that had some conflict... will try to find any memos. Then getting Steve Wolff to run it was fortunate. It was obvious that the centers shouldn't be driving the network.

6) Sometimes people believe that NSF funding is a zero-sum game. This might mean there was resistance from other directorates to the creation and growth of CISE. Did you experience this attitude? How did you deal with it?

In fact CISE was really a tree pruning exercise. No new money came in. Yes. Engineering, MPS also lost funding and organization.

Erich really managed this, but the other Directorate manager (e.g. Rich Nichols I believe) managed this. I had strongly believed in an independent Directorate for some time.

a) There was already a computer engineering program in the engineering directorate, which got moved over into CISE over the objection of the Engineering AD. Would you care to talk about this program?

Nam Suh lost some folks, but this was mostly a hardware and devices. Note semiconductor research stayed in engineering

I believe Nam, I, and Erich all believed that CS is substantially an engineering discipline.

b) Was there any effort to bulk up the size of CISE, so that it really looked like a directorate and so that there would be less incentive for some successor of Eric Bloch as NSF Director to turn CISE into a division in another directorate?

No. CISE was well funded.

7) You had had a heart attack not long before you came to NSF. Did this have any bearing on how you carried out your work at NSF, e.g. the division of labor of work in the CISE office?

I had this 3 years earlier and that caused me to leave Digital and I went to start Encore computer and also to do angel investing. The only affect it may have had was to make me a little more remote i.e. not take the decisions as personal when I forced the decisions that had conflict. The main thing was Erich gave me the support for all the changes

8) Chuck Brownstein has had some very favorable comments to make about your management style. Can you talk about your management style and how effective it you perceived it to be?

In a way, I never perceived myself as a manager, but really more as a leader. I assume that the Division heads were able to manage the NSF process that, left alone, just moves money. I felt my job was to force the changes that were necessary e.g. like closing von Neumann center when it was really not needed and CDC had failed. Also see the attached re. the ASC funding as it evolved

What were the greatest management challenges and successes during your time in CISE?

I tended to be at odds with the centers folks and for them to get more external support to validate their need. See attached.

Also wanted them to focus as a single facility vs. fiefdoms that characterize supers centers.

The biggest success by far was writing the response to the Gore Bill that eventually begot the Internet later

on.

I believe this is the only thing that was ever done across agencies: nsf, darpa, doe, nasa, NIH, doc, etc.

NSF really drove this by the funding of the regional net funding and the document and then getting Bob Kahn's organization to take on the management that got Michigan.

9) I understand that you shared Eric Bloch's desire for larger dollar-value grants and grants of longer duration, and that there was some resistant from the program officers for that meant they could not provide support to as wide a portion of the community. Is this correct?

Yes. The centers grants were going when I arrived. I had been a reviewer of the first centers grants. Also, had experience with DARPA vs NSF small grant funding.

How did you handle this situation?

I don't recall whether we did any of these.

10) Rick Adrion has shared a brief document with me from the time that you prepared for a discussion with the NSF Board and Director. This indicates that your five proposed areas of emphasis for CISE were:

- Parallelism, applied to parallel processing
- Automation, robotics, and intelligent systems
- Ultra-large scale integrated systems
- Advanced scientific and **engineering** computing [emphasis in original]
- Networks and distributed computing

Can you discuss the negotiations that took place as you discussed these programmatic emphases for CISE? Did the programs get modified through this process of negotiation?

Early on, I recall having an advisory panel that reviewed our dircitnn

11) Another document that Rick Adrion has shared with me as an NSF org chart that indicates the 4 proposed organizational structures (divisions) that corresponded to these programmatic initiatives:

- Computer research
- Information science and technology
- Computer and information engineering
- Advanced scientific computing

These evolved into CCR, IRIS, MIPS, and ASC. Later, networking split off from ASC. Can you discuss these decisions about organizational structure?

The formation of a network group that was separate from ASC came from the obvious pruning that I believe I determined to be more important than ASC.

This was something that I took away from the very strong supercomputer centers. This stemmed from a long personal belief and support of networking that started at DEC and continued through the ARPANET.

How well did they achieve your programmatic goals? NREN begot Internet as the result. What were the challenges? Just doing it.

12) Both networking and high-performance computing were important parts of the story of the development of computing at NSF yet somewhat apart from mainstream computer science. Can you talk about their status when you arrived at NSF, your actions in these areas, how these activities changed over time, and how they affected the support of other areas of computer science education and research?

- a) In the networking area: what was the relationship between NSF and DARPA?
- b) Basically DARPA had gotten out of the networking business, but their budget was being eaten in support of ARPAnet that was consumed with email transmission.

With other agencies? See the agencies and people in the FCCSET documents. This was the seminal group, including network group.

# FCCSET COMMITTEE ON COMPUTER RESEARCH AND APPLICATIONS

Paul G. Huray (Chair)  
Office of Science and Technology Policy

## SUBCOMMITTEES

### Science and Engineering Computing

James F. Decker (Chair)  
Department of Energy

James Burrows  
National Bureau of Standards

John S. Cavallini  
Health and Human Services

Melvyn Ciment  
National Science Foundation

John Connolly  
National Science Foundation

Craig Fields  
Defense Advanced Research  
Projects Agency

Harlow Freitag  
Supercomputer Research Center

Randolph Graves  
National Aeronautics and Space  
Administration

Norman H. Kreisman  
Department of Energy

Lewis Lipkin  
National Institutes of Health

Allan T. Mense  
Strategic Defense Initiative Office

David B. Nelson  
Department of Energy

C. E. Oliver  
Air Force Weapons Lab

John P. Riganati  
Supercomputer Research Center

Paul B. Schneck  
Supercomputer Research Center

K. Speierman  
National Security Agency

### Computer Research and Development

Saul Amarel (Chair)  
Defense Advanced Research  
Projects Agency

Donald Austin  
Department of Energy

C. Gordon Bell  
National Science Foundation

James Burrows  
National Bureau of Standards

Bernard Chern  
National Science Foundation

Peter Freeman  
National Science Foundation

Lee Holcomb  
National Aeronautics and Space  
Administration

Charles Holland  
Office of Naval Research

Robert E. Kahn  
Computer Science Technology  
Board

Daniel R. Masys  
National Institutes of Health

Robert Polvado  
Central Intelligence Agency

David Sadoff  
Department of State

William L. Scherlis  
Defense Advanced Research  
Projects Agency

K. Speierman  
National Security Agency

Stephen L. Squires  
Defense Advanced Research  
Projects Agency

Charles F. Stebbins  
Air Force Systems Command

Daniel F. Weiner, II  
Joint Tactical Fusion Program

### Computer Networking, Infrastructure and Digital Communications

C. Gordon Bell (Chair)  
National Science Foundation

Ronald Bailey  
National Aeronautics and Space  
Administration

Sandra Bates  
National Aeronautics and Space  
Administration

James Burrows  
National Bureau of Standards

John S. Cavallini  
Health and Human Services

Thomas Kitchens  
Department of Energy

James Oberthaler  
National Institutes of Health

Dennis G. Perry  
Defense Advanced Research  
Projects Agency

Arnold Pratt  
National Institutes of Health

Shirley Radack  
National Bureau of Standards

Rudi F. Saenger  
Naval Research Laboratory

Daniel VanBelleghem  
National Science Foundation

Stephen Wolff  
National Science Foundation

What role, if any, did Al Gore play here?

Gore wrote the Gore Bill which I believe was helped by the staff.

The essence of the bill (Public Law 99-383, August 21, 1986) is given in the FCCSET Vol. I report.

Is there anything you wish to say about Steve Wolf's role in this area?

Steve was an important figure based on his experience and administrative capability.

c) In the high-performance computing area, how did NSF activities relate to those of DOE and NASA?



- d) tSee the FCCSET structure.
- e)
- f) Do I understand correctly that you already had an interest in high-performance computing from your time at Carnegie Mellon?
- g) This is a long term interest of mine, including the computers that were built at CMU.
- h) I wrote a number of papers on the history, progress and alternatives to get to various performance levels e.g. mega, giga, tera, exa

Is there anything you wish to say about John Connelly's role in this area?

John was a great advocate for supercomputing centers, having established temporary contracts prior to the centers.

I tended to focus on the cost, utilization and effectiveness versus starting more centers.

13) What can you tell me about your interactions with Congress? For example, dealing with Rep. Edward Boland in appropriations? Handling the 10% cut that was caused by Gramm-Rudman balanced budget amendment? Other particular incidents or key figures in Congress or the White House that should be mentioned?

I attended some of the congressional hearings, but never had to present, but may have answered a budget question or two. I came away with enormous respect for the bureaucrats as being much brighter and harder working than the politicians.

14) You only stayed at NSF for a year and a half. Why did you decide to leave?

I left to head up engineering at Dana, a startup I had helped found in Silicon Valley. The head of engineering had been fired.

I felt that CISE was up and running with the right structure and division heads and I wasn't essential and I had completed the formation of CISE.

Re me cutting out prematurely, I think my agreement with Erich was for me to start CISE...in any case, I prefer to startup organizations.

Was that enough time to have an impact?

Yes. The org structure, division heads, and charters/plans were exactly where I thought they should be.

Can you give me a snapshot of the state of CISE at the time you left?

Org and charters and plans were all in place.

Job now was to finish building NREN with NSF and the rest of the government agencies. This had been contracted with CNRI.

Big thing that didn't get done was to integrate computational science with computer science. This is yet to happen.

What were the major opportunities and challenges facing CISE at that time?

ASC funding and effectiveness and coupling with industry ... also expanding user base versus more capacity for QCD computation.

Do you believe the directions that you set for CISE were continued after your departure?

I think Bill Wulf may have told me, that he just implemented the established plan.

15) More broadly, what were the most memorable or key events at NSF and in CISE, e.g., re-organizations, programs, initiatives, etc.?

Just getting the CISE organization, especially networking, in place that functioned as one with external reviews.

The most significant was the NREN Plan responding to the Gore Bill requesting NSF to come up with a plan and proposal for networking that was done with all the other agencies, followed up with the contract to CNRI (Kahn's organization) that subcontracted the building of out of NREN at U of Michigan and IBM.

The Network plan that responded to the Gore Bill then stimulated/embarrassed the two other FCCSET groups (see the figure) to want to leverage a concrete plan to advance their agendas of Decker for HPC at DOE, and Saul Amarel for CS Research at DARPA. Thus our plan became part of an overall plan attached that Steve Squires and Bill Scherlis really got together. This plan morphed into a book then was updated and republished over the next few years as a shiny blue book from FCCSET. Unfortunately, I don't have copies of those books. In the mid to late 90s? CSTB I believe published a very aggressive proposal for CS funding that had everything in it... but went nowhere.

With ASC, I forced the focus on manufacture supplied UNIX software in an effort to make it easy to migrate apps across systems.

In focusing on parallelism, I was motivated to offer the Gordon Bell Prize for parallelism that the ACM administers at the Supercomputing conference. I believe the prize got the whole thing started. See my keynote

[https://www.researchgate.net/publication/316283770\\_Gordon\\_Bell\\_Prize\\_Three\\_Decades\\_Motivating\\_and\\_measuring\\_High\\_Performance\\_Computing\\_progress](https://www.researchgate.net/publication/316283770_Gordon_Bell_Prize_Three_Decades_Motivating_and_measuring_High_Performance_Computing_progress)

Just being downtown, living near Dupont Circle, walking to work in the humid DC and being next to the Whitehouse (where I did a bit of photo bombing) was personally enjoyable. I don't think I would like to be away from the center as NSF now is.

16) I understand that the NRC CSTB had become moribund and that it was revived and supported well by CISE at about the time that you were AD?

I don't recall any interaction with the group.

Later on, I convinced them to look at the HPC systems as a study. I believe Fred Brooks and Ivan Sutherland led this study.

My goal was to try and get focus on parallel programming since it was clear that microprocessor had a dramatic advantage in price-performance

What function, if any, did CSTB play in helping to identify new directions for the field, or legitimizing the creation of new CISE programs?

None that I recall

Was this activity modeled after work that was already being done in the physics community?

No connection that I know of or can recall.

17) Chuck Brownstein was the acting AD both before your arrival and in the interim between your term and that of Bill Wulf. Is there anything you wish to say about Chuck's role as either acting AD or as your deputy?

Chuck was simply critical to help get around and understand the political environment.

Often I would write a memo, Chuck would translate it into NSF Bureaucratise... he also said that my printer needed to be connected to a shredder.

18) Who were the most memorable/influential people with whom you worked/interacted while you were at NSF (NSF and external) who we have not already discussed?

I had good an long lasting interactions with the supercomputer community at DOE that went back to my first visit to Livermore in 1961 or so.

Why were they memorable and/or influential?

Bloch established CISE. It is unclear whether this would have ever occurred without him doing it.

He was supportive a manager and kind of delighted at seeing me in hot water: when I closed von Neumann; had to deal with CDC; had to write a letter to the Pres of Bell Labs when I had been interviewed and said that ATT would only screw up NCR that they had just bought; respond to Cornell that their Nobel Prize winner couldn't spend the money they had been given for running the supercomputer on a crazy research project that we had rejected.

He introduced a 360 evaluation of the organization whereby everyone was peer reviewed. NSF under him was quite open with conflicts exposed.

He was effective at getting a big increase to the budget taken over a few years during the Reagan years.

19) Can you comment on influences from the computer science discipline, industry, the government, or society more generally (if any) that shaped the development of CISE or your programs at the time you were at NSF?

DARPA was in the midst of the Strategic Computing Initiative

20) Any other comments?

I noted that something about the CMU community that felt that NSF and government service was important: me, Bill Wulf, Nico Haberman, Peter Freeman, and Jeannette Wing,