REPORT

OF THE

PANEL ON LARGE SCALE COMPUTING

IN SCIENCE AND ENGINEERING

Peter D. Lax, Chairman

Under the sponsorship of

Department of Defense (DOD)
National Science Foundation (NSF)

In cooperation with

Department of Energy (DOE)
National Aeronautics and Space Administration (NASA)

December 26, 1982

EXECUTIVE SUMMARY

Report of the Panel on Large Scale Computing
in Science and Engineering

Large scale computing is a vital component of science, engineering, and modern technology, especially those branches related to defense, energy, and aerospace. In the 1950's and 1960's the U. S. Government placed high priority on large scale computing. The United States became, and continues to be, the world leader in the use, development, and marketing of "supercomputers," the machines that make large scale computing possible. In the 1970's the U. S. Government slackened its support, while other countries increased theirs. Today there is a distinct danger that the U.S. will fail to take full advantage of this leadership position and make the needed investments to secure it for the future.

Two problems stand out:

> Access. Important segments of the research and defense communities lack effective access to supercomputers; and students are neither familiar with their special capabilities nor trained in their use.

Access to supercomputers is inadequate in all disciplines. Agencies supporting some disciplines such as fusion energy, atmospheric sciences, and aerodynamics have funded National computing facilities through which their remote users have limited networking capabilities. In those disciplines that attempt to fund computing through individual research grants, access to large scale computing remains minimal.

> Future Supercomputers. The capacity of today's supercomputers is several orders of magnitude too small for problems of current urgency in science, engineering, and technology. Nevertheless, the development of supercomputers, as now planned in the U.S., will yield only a small fraction of the capability and capacity thought to be technically achievable in this decade.

Significant new research and development effort is necessary to overcome technological barriers to the creation of a generation of supercomputers that tests these technical limits. Computer manufacturers in the U. S. have neither the financial resources nor the commercial motivation in the present market to undertake the requisite exploratory research and development without partnership with government and universities.

Unless these barriers are overcome, the primacy of U. S. science, engineering, and technology could be threatened relative to that of other countries with national efforts in supercomputer access and development. Although the Federal Government is the first and by far the largest customer for supercomputers, there are no national plans to stimulate the development and use of advanced computer technology in the U. S.

Recommendations:

The Panel recommends the establishment of a National Program to stimulate exploratory development and expanded use of advanced computer technology. The Program has four principal components, each having short- and long-term aspects. Underlying them all is the establishment of a system of effective computer networks that joins government, industrial, and university scientists and engineers. The technology for building networks that allow scientists to share facilities and results is already developed and understood; no time should be lost in connecting existing research groups and computing facilities.

The four components of the recommended program are:

1. Increased access for the scientific and engineering research community through high bandwidth networks to adequate and regularly updated supercomputing facilities and experimental computers;

2. Increased research in computational mathematics, software, and algorithms necessary to the effective and efficient use of supercomputer systems;

3. Training of personnel in scientific and engineering computing; and

4. Research and development basic to the design and implementation of new supercomputer systems of substantially increased capability and capacity, beyond that likely to arise from commercial requirements alone.

The Panel recommends that this program be coordinated within the Federal Government by an interagency policy committee, and that an interdisciplinary Large Scale Computing Advisory Panel be established to assist in its planning, implementation, and operation.

The Panel believes that current funding levels are insufficient to maintain the Nation's leadership in large scale computing. Federal agencies that depend on large scale computing to fulfill their missions must work together to reexamine priorities and to create a coherent program responsive to their individual missions. The Panel has set forth policy and planning issues and has outlined some options for implementation.

Members of the
Panel on Large Scale Computing
in Science and Engineering

William F. Ballhaus, Jr.	NASA-Ames Research Center
James C. Browne	University of Texas, Austin
Brice Carnahan	University of Michigan
Michael Creutz	Brookhaven National Laboratory
Richard H. Gallagher	University of Arizona
Herbert Keller	California Institute of Technology
John Killeen	Lawrence Livermore National Laboratory
Peter D. Lax, Chairman	New York University, National Science Board
Walter M. Macintyre	National Center for Atmospheric Research
Steven Orszag	Massachusetts Institute of Technology
Werner Rheinboldt	University of Pittsburgh
Allan Robinson	Harvard University
Jacob Schwartz	New York University
Edward Wilson	University of California, Berkeley
Kenneth Wilson	Cornell University

CONTENTS

I. Introduction

Supercomputers are the fastest and most powerful scientific computing systems[1] available at any given time: they offer speed and capacity, or special characteristics, significantly greater than on the most widely available machines built primarily for commercial use.[2] Large scale scientific computing is the application of supercomputers to the solution of a model or simulation of a scientific or engineering problem through the appropriate use of numerical algorithms and techniques.

The availability of supercomputers during the past thirty years has been crucial to the Nation's advances in science, engineering, national security, and industrial productivity. Supercomputers have been essential to scientific and engineering investigations in areas such as atmospheric research, astrophysics, molecular biology, integrated circuit design, and fusion research. The weapons programs of DOE, cryptographic analysis, and weather forecasting are dependent on the availability of computational facilities. The use of supercomputers in the aerospace, petroleum, semiconductor, and nuclear industries contributes substantially to the nation's productivity. The development of supercomputers has significant spinoffs for all the technologically based components of the national economy. Research and development in semiconductor technology and in computer research has directly supported and expanded the defense, industrial, medical, and consumer segments of the economy.

The U.S. is the acknowledged leader in the development and use of supercomputers. In 1970 this Nation was preeminent in all aspects of electronic, computer, and computational technology. However, America's present leadership in supercomputers is challenged in the areas of components, development, and applications. Recently, Hitachi has begun marketing what is claimed to be the first 16k bipolar ECL RAM[3]; this device, representative of the continuing advances of Japanese microelectronic manufacturers, is designed for applications in today's scientific computers. Fujitsu, Nippon Electric Company, and Hitachi have each developed supercomputers, which are claimed to compare favorably with the available, or announced, American systems. American universities and research centers, which have historically created new applications of

---

1 These include specialized machines, such as array processors, that are equal to or, for some problems, more powerful than general purpose mainframes.

2 The term supercomputer, as used in this report, encompasses hardware, software, supporting peripherals, and the facilities and personnel needed for their appropriate use. Appendix I reproduces three papers presented to the Panel: NCAR Computing Capabilities and Services, by W. Macintyre; Magnetic Fusion Energy and Computers, by J. Killeen; and The Potential of Los Alamos National Laboratory to Provide Large Scale Computational Capabilities to the Research Community, by B. Buzbee and D. Sparks. These papers present descriptions of supercomputer facilities and of their access from remote facilities through networking.

3 See Electronic Engineering Times, December 6, 1982.

supercomputers and sustained research in computational mathematics, have lagged behind their Japanese and European counterparts in the installation of supercomputers.

Significant national thrusts in supercomputing are being pursued by the governments of Japan, West Germany, France, and Great Britain. Some of these, notably the Japanese effort, center on the development of supercomputers; others, on the provision of supercomputers, or access to them through networks, to the scientific and engineering research community. The British program, for example, is designed to provide research scientists and engineers in academic and government laboratories access to supercomputers through modern workstations connected to a high-speed national network. These agressive foreign national initiatives provide a striking contrast to the current state of planning in the U.S.. The domestic computer industry continues its vigorous research and development efforts in the supercomputer field; however, it is felt that these efforts, necessarily dictated by commercial conditions, are less than they could be and far less than should be for the national scientific and technical capability as a whole. The U.S. research community does not have sufficient access to supercomputing facilities, as is documented in numerous studies, papers, and reports directed toward specific disciplines and specific agencies. A partial bibliography of these studies is included in this report.

Expressions of concern that the U.S. is failing to exploit its position of leadership in supercomputing are being voiced from many quarters. Reflecting this concern, the NSF/DOD Coordinating Committee requested, in April of 1982, that a workshop be organized to explore the problems, needs, and opportunities in large scale computing. This Workshop, sponsored by NSF and DOD with the cooperation of DOE and NASA, was led by a panel of fifteen scientists and engineers from a broad spectrum of disciplines. It took place at the NSF on June 21-22, 1982, and was attended by over one hundred participants. Experts in the use, design, and management of large scale computers from the computing, defense, and other industries, government laboratories, universities, and research funding agencies were included. The lists of the participants in this Workshop are contained in the Supplement.

The Panel assessed the role of supercomputing in scientific and engineering research; surveyed the current use, availability, and adequacy of supercomputers; and considered near- and long-term needs. Subsequent to the June 21-22 Workshop, numerous meetings of smaller groups of participants have taken place; in particular, experts on computer development (Group 3 of the Lax Panel) met at Bellaire, Michigan, on August 23-24, 1982, to further explore avenues for assuring the development of future supercomputers. From these meetings a large number of suggestions and position papers have been directed to the Panel and to the Organizing Committee. This report is an attempt, on the part of the Organizing Committee and the Panel, to outline both the results of the Workshop and the subsequent discussions and contributions. The Panel has chosen not to repeat all the detailed technical arguments or examples of the use of supercomputers found in the literature. A bibliography and appendices are included.

Overall, this report outlines the issues and options for the U.S. to maintain its leadership in supercomputers and supercomputing. Because the issues involve many Federal agencies, government laboratories, universities, private sector companies, and scientific disciplines, they need to be addressed on a National basis and require Federal study, planning, and support. The Panel's report attempts to bring the fundamental issues to the attention of policymakers; however, it deliberately avoids details of an organizational, programmatic, or budgetary nature.

II Summary of Findings and Recommendations

Summary of Findings

Large scale computing is a vital component of science, engineering, and technology[4], bringing together theory and applications. It is essential for the design of many technologically sophisticated products, and is making possible for the first time the analysis of very complex scientific and engineering problems which to date have defied analytical and experimental techniques. Examples of the importance of supercomputing are briefly noted below[5].

Renormalization group techniques[6] are a major theoretical breakthrough that provide a new framework for the understanding of a number of unsolved scientific and engineering problems ranging from problems in quantum field theory, the onset of phase transitions in materials, the development of turbulence, propagation of cracks in metals, and the exploitation of oil reservoirs. Only a minute fraction of these problems can be solved analytically. Large scale computational techniques have been essential to the use of renormalization group methods, and even today's largest computational machines are not sufficiently powerful to address most of these problems[7].

Aerodynamic design using a supercomputer has resulted in the design of an airfoil with 40% less drag than that developed by previous experimental techniques[8]. The solution of the Navier-Stokes equations with sufficient

------

4 See, for example, R & D for National Strength, Center for Strategic and International Studies, Georgetown University, Washington, D.C., 1982. See also The Defense Science Board Summer Study, 1981.

5 In the Supplement of this Report, H. B. Keller and J. R. Rice describe in some detail scientific and engineering areas in need of supercomputers. Appendix II contains a number of examples of scientific and engineering problems successfully addressed on supercomputers as well as additional examples requiring supercomputing capabilities not yet available.

6 These techniques were devised to handle the description of phenomena with interactions spanning an extremely wide scale. K. G. Wilson was awarded the 1982 Nobel Prize in Physics for his contributions to the theoretical development and application of renormalization group techniques to critical phenomena. See his contributions to this report in the Supplement, as well as in Appendices II and III.

7 See, in Appendix II, several problems posed by K. G. Wilson.

8 For a detailed description see "Trends and Pacing Items in Computational Aerodynamics", [40].

resolution to represent faithfully fluid behavior became practical with the current (Class VI) generation of supercomputers. The wings of the Boeing 767 and of the European Airbus 310 were designed by computational methods on such machines, resulting in this most significant improvement.

The aerodynamic design of an entire aircraft is not feasible with today's supercomputers; it is estimated that machines 100 times more powerful are needed for this purpose. The design of jet engines, involving the simulation of complex three-dimensional fluid flows and associated chemical reactions, also requires significantly increased computational capability and capacity.

In design, especially of advanced weapons systems, large scale computational modeling is an essential substitute for experimentation. Similarly, the design of future generations of nuclear power plants, and their operation—relying on real-time simulation for their control—require computational facilities several orders of magnitude greater than those available today.

Perhaps the most significant applications of scientific computing lie not in the solution of old problems but in the discovery of new phenomena through numerical experimentation; the discovery of nonergodic behavior, such as the formation of solitons, and the presence of strange attractors and universal features common to a large class of nonlinear systems are examples of this scientific process.

Current and feasible supercomputers are extremely powerful scientific and engineering tools. They permit the solution of previously intractable problems, and motivate scientists and engineers to explore and formulate new areas of investigation. They will surely find significant applications not yet imagined. For these reasons, the Panel believes that it is in the National interest that access to constantly updated supercomputing facilities be provided to scientific and engineering researchers, and that a large and imaginative user community be trained in their uses and capabilities.

The U.S. has been and continues to be the leader in supercomputer technology and in the use of supercomputers in science and engineering. The present position of leadership is evidenced by the dominance of the supercomputer market by American producers and by the successful exploitation of supercomputing at national laboratories. However, the Panel finds that this position of leadership is seriously undermined by the lack of broad scale exploration, outside of a few national laboratories, of the scientific and engineering opportunities offered by supercomputing, and by a slowdown in the introduction of new generations of supercomputers. This threat becomes real in light of the major thrust in advanced supercomputer design that is being mounted by the Japanese Government and industry, and by vigorous governmental programs in the United Kingdom, West Germany, France, and Japan to make supercomputers available and easily accessible to their research and technological communities.

American preeminence in large scale computing has been a result of the confluence of three factors: the vitality of the U.S. computer industry, the far-sighted policies of the Federal government, and the leadership of scientists and engineers from universities and government laboratories. The Atomic Energy Commission, on the urging of John von Neumann, initiated the use of large scale computation in research and weapons design; NASA, prodded by Hans Mark, advanced the use of supercomputing in its scientific programs. American universities and government laboratories conducted the research that formed the basis for constructing and applying computers, trained the needed scientific and engineering personnel, and made computers and computing an essential tool in scientific and engineering research.

The Federal government vigorously implemented policies that supported these efforts, granted generous funds for computation, and, through its role as the major purchaser of scientific computers, provided the incentives and insured the market for these unique machines. Forward-looking corporations exploited the scientific and engineering opportunities, developed an advanced industrial technology, and created this most vital component of the American economy.

During the 1970's the Federal government retreated from its support of large scale computing in universities. The NSF program to provide and expand university computing facilities for scientific and engineering research was terminated in 1972; at about the same time IBM discontinued its generous discounts for the purchase of computing equipment by academic institutions as a result of pressures from the Justice Department and competitors. Since then large scale university computing facilities have withered while the action shifted to national laboratories and to industrial users. The most advanced scientific computer of the early seventies, the CDC 7600, was not installed on a single American campus, although it was available to researchers at several foreign universities and research institutes.

This continues today. With the exception of two universities and a few government laboratories, either dedicated to special tasks or specific disciplines, universities and government research installations lack supercomputers.[9]

Within the Government, fully integrated supercomputer facilities are found exclusively at dedicated national laboratories such as Los Alamos National Laboratory (LANL) and Lawrence Livermore National Laboratory (LLNL) in support of weapons and fusion programs; and NASA installations, the Geophysical Fluid Dynamics Laboratory, and the National Center for Atmospheric research (NCAR) in support of aerospace, oceanographic, and atmospheric research programs. The National Magnetic Fusion Energy Computer Center (NMFECC) of LLNL is accessible in interactive mode to researchers at remote locations by a high speed network.[10] On the other hand, the National Bureau of Standards and most DOD laboratories do not have supercomputers and far too few universities have the specialized computational equipment needed for scientific computing (e.g., array processors). As a result of limited access to supercomputing facilities by the broad research community, significant research opportunities[11] have

9 See, in Appendix I, Partial Inventory and Announced Orders of Class VI Machines.

10 See Appendix I for a description of some of these facilities.

11 See, for example, the "Prospectus for Computational Physics," [2], "Future Trends in Condensed Matter Theory and the Role and Support of Computing" [4], "Report by the Subcommittee on Computational Capabilities for Nuclear Theory," [28], and "An Assessment of Computational Resources Required for Ocean Circulation Modeling," [35].

been missed, and the younger generation of researchers is inadequately trained in large scale computing.

The need for access to large scale computational facilities has become so critical[12] that several universities, assuming significant financial risks, have felt it essential to acquire supercomputers. Several more are seriously considering doing so, and others are in the process of forming consortia for this purpose. Some of these endeavors have applied for Federal funding without which they may have financial difficulties. Other groups are pressing funding agencies to expand or replicate highly successful facilities, such as those at NCAR, NMFECC, and NASA-Ames, at universities or at national laboratories. Class VI scientific remote computing services are available from a few commercial service bureaus, but neither the academic nor the government research communities make extensive use of this resource. This seems due to a combination of lack of funds for computing services, the perceived high cost associated with these services, and a lack of sophisticated high-speed networking facilities. It is an indication of the absence of a national plan that a substantial number of leading scientists are clamoring for access to supercomputers at the same time that some supercomputing facilities are underutilized.[13]

A supercomputer is a general purpose scientific instrument serving a broad and diverse base of users. The decline of supercomputing at universities is analogous to the decline of instrumentation; neither large scale computing nor instrumentation can be sustained in a stable manner through funding of individual research grants, where their costs must compete with that of scientific personnel.[14]

The findings of the Panel regarding the development of supercomputers are as alarming as the findings on their access and availability. The U.S. supercomputing market is, at this time, dominated by Cray Research (CRAY-1) and Control Data Corporation (CYBER 205).[15] The Japanese vendors, Hitachi (S-210/20) and Fujitsu (VP-200), have announced the delivery of supercomputers in the near future, and these machines appear to be comparable to the available American systems. The Japanese are striving to become serious competitors of domestic manufacturers, and U.S. dominance of the supercomputer market may soon be a thing of the past. The Japanese Government-sponsored National Super Computer Project[16] is aimed at the development, by 1989, of a machine one thousand times faster than current machines. There is no comparable technical program in the U.S.. The Panel notes that in the case of the NASA Numerical Aerodynamic Simulator, a very high performance supercomputer, no acceptable proposals for its development

---

12 See, in Appendix III, The Supercomputer Famine in American Universities, by L. L. Smarr.

13 Supercomputer cycles are available at the University of Minnesota CRAY-1 and at Colorado State University CYBER 205.

14 See, in Appendix III, the paper by R. G. Gillespie. Most supercomputer facilities are funded directly by the Federal Government.

15 See, in Appendix I, Partial Inventory and Announced Orders of Class VI Machines.

16 See, in Appendix III, Japan's Initiatives in Large Scale Computing, by L. Lee.

were received. Neither of the two competing vendors could assure NASA that they would meet the performance requirements. Rather than developing new products, the vendors attempted to fit all NASA requirements to their existing product line.

Upon review of previous studies, the Panel also finds that the power of current and projected supercomputers is insufficient to meet existing needs in science, engineering, and technology, both military and civilian.[17] Research at universities and in the computer industry has indicated that future generations of very high performance computer systems may have parallel architectures radically different from the conceptually sequential architectures of today's supercomputers. There are many candidate architectures that must be evaluated before commercial feasibility can be established.[18] Simultaneously, the rapid and continuing advance of microelectronic technology makes it feasible to build such parallel machines. There is also a need for improvement of component performance.

The Panel believes that under current conditions there is little likelihood that the U.S. will lead in the development and application of this new generation of machines.[19] Factors inhibiting the necessary research and advanced development are the length and expense of the development cycle for a new computer architecture, and the uncertainty of the market place. Very high performance computing is a case where maximizing short-term return on capital does not reflect the national security or the long-term national economic interest. The Japanese thrust in this area, through its public funding, acknowledges this reality.

The Panel estimates[20] that the present annual investment in basic research on algorithms, software, and architecture is between 5 and 10 million dollars, while the annual advanced development expenditures for supercomputers (beyond Class 6 machines) are between 20 and 40 million dollars. This is contrasted with the development cost for a new high-speed conventional architecture system of approximately 150 million dollars, as well as the estimated 200 million dollars national superspeed computer project in Japan. The panel considers current levels of United States investments insufficient to maintain leadership in supercomputers.

The Panel believes that U.S. leadership in supercomputing is crucial for the advancement of science and technology, and therefore, for economic and national security.

---

17 See, for some illustrative examples, Appendix II. Also notable is "Trends and Pacing Items in Computational Aerodynamics", by D.R. Chapman, [40].

18 See, in Appendix III, a series of contributions to the Workshop. In particular, the papers by Dennis, Gajski, et al., Ris, and Fernbach; also the report of Group 3 of the Lax Panel in the Supplement.

19 See, in Appendix III, Why the U.S. Government Should Support Research on and Development of Supercomputers, by B. Buzbee.

20 Members of the Panel and of the Organizing Committee have conducted a survey to estimate the current total national investment in research and advanced development for supercomputer and supercomputing, both public and private. This survey included research and development costs but excluded funding for the acquisition, maintenance, and operation of supercomputer facilities.

Recommendations

The Panel recommends that the present needs and challenges to U. S. leadership in scientific computing and supercomputer technology be addressed with high priority. To this end, the Panel has set forth the background for planning and policy issues, outlined some options, and noted that current total funding in this area is insufficient to maintain the Nation's leadership in large scale computing. The Panel has avoided recommendations of a programmatic and organizational nature; these, and their implementation, are best left to the appropriate government agencies. These agencies must work together to respond to the issues raised and put together a detailed coherent program whose components are responsive to their individual missions. The program plan should contain a clear statement of goals, directions, and roles for the academic, industrial, and Federal government segments; responsibilities of the participating Federal agencies; and funding required.

The Panel recommends that a long-term National Program on Large Scale Computing should be initiated immediately, with the participation of the appropriate Federal agencies, the universities, and industry. The goals of this National Program should be:

1. Increased access for the scientific and engineering research community through high bandwidth networks to adequate and regularly updated supercomputing facilities and experimental computers;

2. Increased research in computational mathematics, software, and algorithms necessary to the effective and efficient use of supercomputer systems;

3. Training of personnel in scientific and engineering computing; and

4. Research and development basic to the design and implementation of new supercomputer systems of substantially increased capability and capacity beyond that likely to arise from commercial sources.

This Program should be coordinated by an interagency policy committee consisting of representatives of the appropriate Federal agencies, including DOC, DOD, DOE, NASA, and NSF. A Large Scale Computing Advisory Panel, with representatives from Government, the universities, and industry, should be established to assist in the planning, implementation, and operation of the Program.

The Panel finds two points that require emphasis:

As the few successful facilities amply demonstrate, solution of the problem of access to supercomputing facilities, on a national basis, is possible.

Recommendations

The Panel recommends that the present needs and challenges to U. S. leadership in scientific computing and supercomputer technology be addressed with high priority. To this end, the Panel has set forth the background for planning and policy issues, outlined some options, and noted that current total funding in this area is insufficient to maintain the Nation's leadership in large scale computing. The Panel has avoided recommendations of a programmatic and organizational nature; these, and their implementation, are best left to the appropriate government agencies. These agencies must work together to respond to the issues raised and put together a detailed coherent program whose components are responsive to their individual missions. The program plan should contain a clear statement of goals, directions, and roles for the academic, industrial, and Federal government segments; responsibilities of the participating Federal agencies; and funding required.

The Panel recommends that a long-term National Program on Large Scale Computing should be initiated immediately, with the participation of the appropriate Federal agencies, the universities, and industry. The goals of this National Program should be:

1. Increased access for the scientific and engineering research community through high bandwidth networks to adequate and regularly updated supercomputing facilities and experimental computers;

2. Increased research in computational mathematics, software, and algorithms necessary to the effective and efficient use of supercomputer systems;

3. Training of personnel in scientific and engineering computing; and

4. Research and development basic to the design and implementation of new supercomputer systems of substantially increased capability and capacity beyond that likely to arise from commercial sources.

This Program should be coordinated by an interagency policy committee consisting of representatives of the appropriate Federal agencies, including DOC, DOD, DOE, NASA, and NSF. A Large Scale Computing Advisory Panel, with representatives from Government, the universities, and industry, should be established to assist in the planning, implementation, and operation of the Program.

The Panel finds two points that require emphasis:

As the few successful facilities amply demonstrate, solution of the problem of access to supercomputing facilities, on a national basis, is possible.

Secondly, the domestic computer industry must allocate its scarce research and development funds to meet all the commercial opportunities and competitive challenges. Supercomputing enjoys a priority within the computer industry. But this priority, which reflects competitive commercial conditions, does not reflect the entire national scientific and security interest. It is not reasonable to rely solely on industry's own initiatives and resources in this area.

Possible Approaches for the National Program

The Panel has received many suggestions[21] for carrying out the thrusts of the proposed National Program. We outline here those considered most promising.

1. Access: There appear to be three approaches to provide reliable and efficient access to supercomputers to the research and development community. Common to all these is the development of a nation-wide interdisciplinary network[22] through which users will have access to facilities. This network should connnect all supercomputer facilities (except those dedicated to very special tasks), including commercial supercomputing centers and experimental machines.

   - The most expedient and perhaps least expensive way to provide supercomputer access to the broad range of scientific and engineering researchers is to enhance supercomputer capacity and staff at existing centers which have demonstrated sophisticated capabilities for providing large scale computing.

   - Provide supercomputers to selected government laboratories without such facilities and make them available to the broad research and development community through networking. In addition, there should be sharing and enhancement of current supercomputer facilities located at universities and government laboratories.

   - Establish additional regional centers at selected universities, interconnected with existing facilities at other universities and government laboratories.

   The existence of a national network would permit combinations of these nonexclusive options, as well as the appropriate use of commercial services. The mechanisms for funding these facilities and allocating access should be carefully studied.[23]

---

21 see position papers in Appendix III.

22 The NMFECC network, described in Appendix I, is the example repeatedly mentioned for emulation because of its high bandwidth. The ARPANET network is often mentioned because of its interdisciplinary nature.

23 See the position paper by R. G. Gillespie in Appendix III.

The above recommendations are related to the problems of access to general purpose supercomputer facilities. It should be noted, however, that there are scientific and engineering problems that can perhaps be better and more economically attacked by specialized supercomputing facilities and by sophisticated array processors.[24] The Panel recommends that, as part of the National Program, significant emphasis be placed on providing this specialized equipment to the research community. Finding the proper balance between investments on these two types of facilities requires a careful analysis, at the multidisciplinary and interagency level.

2. Research in Software and Algorithms

Today's supercomputers are a major departure from traditional sequential machines. Future significant improvements may have to come from architectures embodying parallel processing elements - perhaps several thousands of processors. In order to exploit today's vector processors and future parallel processors, entirely new algorithms must be conceived. Research in languages, algorithms, and numerical analysis will be crucial in learning to exploit these new architectures fully. The contributions of numerical analysis, computational mathematics, and algorithm design to the practice of large scale computing is as important as the development of a new generation machines.

3. Training

Another important component of this National Program is the development of an imaginative and skilled user community in supercomputing. There is a considerable shortage of appropriately trained personnel and of training opportunities in this area. Forms of institutional encouragement, such as NASA's special fellowships in the area of numerical fluid mechanics, special summer schools, and special allocation of access time to supercomputers for those projects that involve graduate students, should be considered. Some of the more mathematical aspects of these activities can be accomplished independently of the machines

----

24 Some of these problems arise in a number of areas associated with experimental physics. See, in Appendix III, the letter from A. E. Brenner, Jr. See, also in the same Appendix, the position paper by K. Wilson on the use of array processors.

on which the actual calculations are done; however, the true integration of methods and their implementation cannot be done without access to supercomputers. The nature of the machine architecture has a very profound effect on the numerical methods, on the algorithms, and of course on the software. Thus, while being trained, students must have access to state-of-the-art computers. Today such training is virtually nonexistent; yet the skills gained from such training are essential to science and engineering.

4. Research and Development for New Supercomputers

There are serious component and architectural problems that must be solved as part of the development of future generations of supercomputers. The unique strengths of industry, universities, and Government laboratories should be brought together for this purpose. A group of panelists from this workshop, with the aid of a number of experts from industry and universities, has produced a report[25] which describes one such program.

Since a great deal of careful analysis and detailed planning is required before the proposed National Program can be implemented, the Panel urges that its recommendations be acted upon as soon as possible.

---

25 See, in the Supplement, A Program for Development of Very High Performing Computer Systems, by J. C. Browne and J. T. Schwartz.

Selected Bibliography

[1] "A Push for Major Improvements in Both Large-Scale Computing Capability and Simulation Techniques for Physical Processes," memo to Users of Large Scale Computing for Simulation in Government, Industry and Universities from Kenneth Wilson, James A. Weeks Professor of Physical Science, Cornell University, July 15, 1981.

[2] "Prospectus for Computational Physics," Report by the Subcommittee on Computational Facilities for Theoretical Research to the Advisory Committee for Physics, Division of Physics, National Science Foundation, March 15, 1981, William H. Press, Chairman.

[3] "A Proposal Submitted to the United States Department of Energy for a National Nuclear Computer Center at the University of Maryland," and Appendix A., "A Report on Computer Needs for Nuclear Theory," William MacDonald, Chairman, July 15, 1979.

[4] "Future Trends in Condensed Matter Theory and the Role and Support of Computation," report by the Condensed Matter Theory Panel, published under the auspices of the council on Materials Science for the Division of Materials Sciences, U.S. Department of Energy, J. D. Joannopoulos and A. N. Berker, Co-Chairman, September 22, 1981.

[5] "A Preliminary Proposal to the Department of Energy submitted by Colorado State University for the Support of a Class VI Computer," by Robert W. Hayman and Gerold R. Johnson, Colorado State University, Fort Collins, Colorado, August 24, 1981.

[6] "NAS Project Status Report for the NAS User Steering Group," prepared by the Numerical Aerodynamic Simulator Project Office, National Aeronautics and Space Administration, Ames Research Center, July 1, 1980.

[7] "The Numerical Aerodynamic Simulator - Description, Status, and Funding Option;" a report prepared at the request of the Committee on Commerce, Science, and Transportation, United States Senate; Dec. 1981, NASA.

[8] Numerical Aerodynamic Simulator; Background, Requirements, and Justification; NASA, Ames Research Center, Revised January 1982.

[9] "Proposal for Research on Parallel Applicative Programming for Multiprocessor Systems," Rodney R. Oldehoeft, Stephen J. Allan, Colorado State University, and Burton J. Smith, Denelcor, Inc., October 1981.

[10] "Development Program, Scientific Computation Center," Florida State University Computing Center, Tallahassee, Florida.

[11] "Second Generation S-1 Systems: Computing and Memory Capacity for Demanding Navy Applications," A Review Presentation for Members of the Navy Technical Community, Phys. Brief 81-103, Lowell Wood, et al., S-1 Project, Lawrence Livermore National Laboratory, January 16, 1980.

[12] "Simulation of Turbulent Shear Flows: What can we do and what have we learned?" presented by invitation at the Seventh Biennial Symposium on Turbulence, Rolla, MO, by W. C. Reynolds, Department of Mechanical Engineering, Stanford University, September 21, 1981.

[13] "A Commentary on Computational Mechanics," J. T. Oden and K. J. Bathe, Applied Mechanics Reviews 31 (1978) pp. 1053-1058.

[14] "Draft Proposal to the National Science Foundation for support of Fellowship Program to Institute for Computation, the Center for Advanced Technology at Colorado State University," George G. Olson, June, 1981.

[15] "Report on the Utilization of Large Scale Computers," Committee on the Utilization of Large Scale Computers in ASGA, Bennett Miller, chairman, July 1, 1977.

[16] "Computing Division Two Year Operational Plan, FY 1981-1982," Robert H. Ewald, W. Jack Worlton, Margery McCormick, Los Alamos National Laboratory, LA-8725-MS, February 1981.

[17] Letter to Dr. Alvin W. Trivelpiece, Director, Office of Energy Research from Dr. David A. Shirley, Director, Lawrence Berkeley Laboratory with Conceptual Proposal for Energy Research Network, September 14, 1981.

[18] "Nuclear Weapons and Computers, Summary: The Role of Large Scale Computers in the DOE Nuclear Weapons Program," the Gold Report, prepared for the Director of Military Applications, U. S. Department of Energy, March 1978.

[19] "DOE Research in Utilization of High-Performance Computers," B.L. Buzbee, W. Jack Worlton, G.A. Michael, G. Rodrigue, LA-8609-MS, Los Alamos National Laboratory, December 1980.

[20] "A Philosophy of Supercomputing," Jack Worlton, LA-8849-MS, Los Alamos National Laboratory, June 1981.

[21] "Department of Energy FY 1983-1987 ADP Long Range Plan," Office of ADP Management, December 1981.

[22] "An Assessment of Trends in the Solid State Sciences," National Research Council, July 1980.

[23] "Future of Nuclear Science," by Ad Hoc Panel on the Future of Nuclear Science, 1977.

[24] "Long Range Plan for Nuclear Science," excerpts from a report by the Nuclear Science Advisory Committee, July 30-August 4, 1979.

[25] "Draft II - Astronomy Survey Committee: Report of the Panel on Data Processing and Computational Facilities," May 19, 1980.

[26] "Workshop on Computational Astrophysics," NASA Ames Research Center, February 1980.

[27] "Magnetic Fusion Energy and Computers: The Role of Computing in Magnetic Fusion Energy Research and Development," Office of Fusion Energy, DOE, October 1979.

[28] "Report by the Subcommittee on Computational Capabilities for Nuclear Theory to the DOE/NSF Nuclear Science Advisory Committee," December 1981.

[29] "Purdue Center for Parallel and Vector Computing," Dennis Gannon, John Rice, Saul Rosen, April 23, 1982.

[30] "Universities and the Future of Computer Technology Development," Kenneth G. Wilson, Laboratory of Nuclear Studies, Cornell University, Ithaca, New York.

[31] "Present vs. Future Supercomputers," notes by Kenneth Wilson, Newman Laboratory, Cornell University, May 1982.

[32] "A Data Flow Supercomputer," Jack B. Dennis, G-R Gao, Kenneth W. Todd, Computation Structures Group Memo 213, MIT Laboratory for Computer Science, Cambridge, Massachusetts, March 1982.

[33] "Foreign Travel Trip Report: Visits with Japanese Computer Manufacturers," R. Ewald, Bill Buzbee, J. Worlton, CDO/82-6782A, Computation Division, Los Alamos National Laboratory, April 5, 1982.

[34] "Supercomputers," Ronald D. Levine, Scientific American, January 1982.

[35] "An Assessment of Computational Resources Required for Ocean Circulation Modeling," Ad Hoc Committee on Computing Resources and Facilities for Ocean Circulation Modeling, Ocean Sciences Board, Commission on Physical Sciences, Mathematics and Resources, National Research Council of the National Academy of Sciences, 1982.

[36] "Nearest Neighbor Concurrent Processor," E. Brooks, G. Fox, R. Gupta, O. Martin, S. Otto, E. De Benedictis; California Institute of Technology, DOE Research and Development Report, CALT-68-867, September 1981.

[37] "Why the U.S. Government Should Support Research On And Development Of Supercomputers," Bill Buzbee, Los Alamos Report LA-UR-82-1968, July 1982.

[38] Los Alamos SCIENCE, Los Alamos National Laboratory, Vol. 2, No. 2 (1981).

[39] Justification for Class VI Scientific Computing Capability at Sandia Laboratories, Albuquerque; Sandia Laboratories, February 1979.

[40] "Trends and Pacing Items in Computational Aerodynamics," D. R. Chapman, Lecture notes in Physics, 141, Springer-Verlag, 1981.

Organization of the Workshop on Large Scale Computing for Science and Engineering

Members of the Panel on Large Scale Computing in Science and Engineering

William F. Ballhaus	NASA-Ames
James C. Browne	University of Texas, Austin
Brice Carnahan	University of Michigan
Michael Creutz	Brookhaven National Laboratory
Richard H. Gallagher	University of Arizona
Herbert Keller	California Institute of Technology
John Killeen	Lawrence Livermore National Laboratory
Peter D. Lax, Chairman	New York University, National Science Board
Walter M. Macintyre	National Center for Atmospheric Research
Steven Orszag	Massachusetts Institute of Technology
Werner Rheinboldt	University of Pittsburgh
Allan Robinson	Harvard University
Jacob Schwartz	New York University
Edward Wilson	University of California, Berkeley
Kenneth Wilson	Cornell University

Organizing Committee of the Workshop on Large Scale Computing in Science and Engineering

Donald Austin	Department of Energy
Melvyn Ciment	National Bureau of Standards
Kent Curtis	National Science Foundation
David Fox	Air Force Office of Scientific Research
J. M. Greenberg	National Science Foundation
Carl Hall	National Science Foundation
E. F. Infante	National Science Foundation
Peter D. Lax, Chairman	New York University, National Science Board
Marshall Lih	National Science Foundation
Lawrence Lee	National Science Foundation
James Pool	Office of Naval Research
Paul Schneck	National Aeronautics and Space Administration

June 14, 1982

AGENDA FOR WORKSHOP ON LARGE-SCALE COMPUTING FOR SCIENCE AND ENGINEERING

Location

National Science Foundation
1800 G Street, N.W., Rm. 540
Washington, D.C. 20550

June 21, 1982

Time	Topic	Leader
8:30 a.m.	Preamble	Lax
9:00 a.m.	Large-Scale Computing Needs	Keller, Robinson
10:00 a.m.	BREAK	
10:15 a.m.	Arrangements for Making Computing Power Available	Orszag
	(Ballhaus, Hayes, Killeen, Macintyre)	
12:00	LUNCH	
1:00 p.m.	Computer Technology in the U.S. (Presentations)	Schwartz
	(Ballhaus, Buzbee, Michael, Schneck, Wilson)	
3:00 p.m.	BREAK	
3:15 p.m.	Discussion (Fernbach, Patton, Olson, Rice)	Lax
6:00 p.m.	Adjournment (except for Panel Members)	

June 22, 1982

Time	Topic	Leader
8:30 a.m.	Charge to Panel	Lax
9:00 a.m.	Group Working Sessions	
12:00	LUNCH	
1:30 p.m.	Plenary Session I	
2:30 p.m.	BREAK	
2:45 p.m.	Group Working Session	
5:00 p.m.	Closing Remarks	Lax

List of Participants of the
Workshop on Large-Scale Computing
in Science and Engineering
Washington, DC
June 21-22, 1982

W. Richards Adrion - National Science Foundation
Carl Amthor - Associated Universities, Inc.
Lloyd Armstrong - National Science Foundation
Donald Austin - Department of Energy
Steven Bachman - Control Data Corporation
William F. Ballhaus - National Aeronautics and Space Administration, Ames
Bruce H. Barnes - National Science Foundation
F. Brett Berlin - Cray Research Incorporated
William Bernard - National Science Foundation
Alfred Brenner - Fermi Laboratory
James C. Browne - University of Texas, Austin
Alphonse Buccino - National Science Foundation
Bill Buzbee - Los Alamos National Laboratory
Andrew Callegari - Exxon Research and Engineering Company
P. Cannon - Star Technology
Jagdish Chandra - Army Research Office
Alan Charlesworth - Floating Point Systems, Incorporated
John C. Cherniavsky - National Science Foundation
Y. T. Chien - National Science Foundation
Melvyn Ciment - National Bureau of Standards
Andrew Collier - Time Magazine
Curtis Collins - National Science Foundation
Michael Creutz - Brookhaven National Laboratory
Kent K. Curtis - National Science Foundation
Carl G. Davis - BMD Advanced Technology Center
Dorothy Deringer - National Science Foundation
A. Erisman - Boeing Computer Service
Sidney Fernbach - Consultant, Information Systems
Judith Flippen-Anderson - Naval Research Laboratory
David Fox - Air Force Office of Scientific Research
J. Gaylor - Boeing Computing Services
Robert Gillespie - University of Washington
Paul Gilmore - Goodyear Aerospace Corporation
Randolph Graves, Jr. - National Aeronautics and Space Administration
James M. Greenberg - National Science Foundation
Lloyd Griffiths - University of Colorado
Carl Hall - National Science Foundation
W. F. Hanrahan - CBEMA
Allen J. Hansen - Burroughs Corporation
Richard Harris - Cray Research Incorporated
W. C. Harris - National Science Foundation
Edward Hayes - National Science Foundation
Kenneth Hoffman - Massachusetts Institute of Technology
E. F. Infante - National Science Foundation
Glenn Ingram - National Bureau of Standards

R. F. Johnson - Floating Point Systems, Incorporated
Ronald E. Kagarise - National Science Foundation
Robert Kahn - Advanced Research Projects Agency
James S. Kane - Office of Energy Research, Department of Energy
Herbert Keller - California Institute of Technology
John Killeen - Lawrence Livermore National Laboratory
Ronald Larsen - National Aeronautics and Space Administration
Peter Lax - New York University - National Science Board
Lawrence Lee - National Science Foundation
Stephen F. Lundstrom - Burroughs Corporation
Walter M. Macintyre - National Center for Atmospheric Research
J. V. Martinez - Department of Energy
Paul Mathews - Office of Military Applications, DOE
George Michael - Lawrence Livermore National Laboratory
K. Mohan - Office of Management and Budget
Thomas Nash - Fermi Laboratory
David B. Nelson - Department of Energy
Robert Numrich - Control Data Corporation
C. E. Oliver - Air Force Office of Scientific Research
George Olson - Colorado State University
Barbara Onestak - National Science Foundation
Steven Orszag - Massachusetts Institute of Technology
A. R. Osborne - Johns Hopkins University
Seymour Parter - University of Wisconsin
Peter Patton - University of Minnesota
Howard Penn - National Science Foundation
Victor Peterson - National Aeronautics and Space Administration, Ames
A. H. Pettifor - Office of Science and Technology Policy
Elliot Pierce - Department of Energy
Irwin Pikus - National Science Foundation
James Pool - Office of Naval Research
Jesse Poore - Georgia Institute of Technology
B. Prosswimmer - Boeing Computing Services
Lawrence Rapagnani - Air Force Weapons Laboratory
Werner Rheinboldt - University of Pittsburgh
John Rice - Purdue University
Fred Ris - International Business Machines
Allan Robinson - Harvard University
Paul N. Romani - National Science Foundation
William G. Rosen - National Science Foundation
Sherwin Ruben - Goodyear Aerospace
John Ryff - National Science Foundation
Russ Saunders - Denelcor, Incorporated
Paul Schneck - National Aeronautics and Space Administration
Jacob Schwartz - New York University
Burton J. Smith - Denelcor, Incorporated
W. R. Smith - Cray Research Incorporated
P. Sorianic, Jr. - Star Technology
Dale Sparks - Los Alamos National Laboratory
Paul G. Spyropoulos - National Science Foundation
Leonard Stans - Air Force Weapons Laboratory
Martin T. Storma - Cray Research Incorporated
Jane Stutsman - National Science Foundation
Alvin Thaler - National Science Foundation

Eric Thompson - National Science Foundation
D. Triantaphyllou - Mitre Corporation
Jean M. Vocar - Goodyear Aerospace
Kenneth Wallgren - National Aeronautics and Space Administration
Hugh Walsh - International Business Machines
Frederick Warther - Denelcor, Incorporated
Edward Wilson - University of California, Berkeley
Kenneth Wilson - Cornell University
Robert Woods - Department of Energy

List of Participants of the
Follow-up Workshop on Large Scale Computing for
Science and Engineering
Shanty Creek, Michigan
August 23-24, 1982

Jim Browne - University of Texas
Kent Curtis - National Science Foundation
Jack Dennis - Massachusetts Institute of Technology
David Fox - Air Force Office of Scientific Research
David Kuck - University of Illinois, Urbana
George Michael - Lawrence Livermore National Laboratory
Ed Oliver - Air Force Office of Scientific Research
Fred Ris - IBM Research
Paul Schneck - NASA/Goddard Space Flight Center
Jack Schwartz - New York University
Burton Smith - Denelcor, Inc.
Larry Snyder - Purdue University
Marty Storma - Cray Research
Bob Voigt - NASA/Langley Research Center
Ken Wallgren - NASA/Goddard Space Flight Center

EXECUTIVE SUMMARY

Report of the Panel on Large Scale Computing in Science and Engineering

Large scale computing is a vital component of science, engineering, and modern technology, especially those branches related to defense, energy, and aerospace. In the 1950's and 1960's the U. S. Government placed high priority on large scale computing. The United States became, and continues to be, the world leader in the use, development, and marketing of "supercomputers," the machines that make large scale computing possible. In the 1970's the U. S. Government slackened its support, while other countries increased theirs. Today there is a distinct danger that the U.S. will fail to take full advantage of this leadership position and make the needed investments to secure it for the future.

Two problems stand out:

> Access. Important segments of the research and defense communities lack effective access to supercomputers; and students are neither familiar with their special capabilities nor trained in their use.

Access to supercomputers is inadequate in all disciplines. Agencies supporting some disciplines such as fusion energy, atmospheric sciences, and aerodynamics have funded National computing facilities through which their remote users have limited networking capabilities. In those disciplines that attempt to fund computing through individual research grants, access to large scale computing remains minimal.

> Future Supercomputers. The capacity of today's supercomputers is several orders of magnitude too small for problems of current urgency in science, engineering, and technology. Nevertheless, the development of supercomputers, as now planned in the U.S., will yield only a small fraction of the capability and capacity thought to be technically achievable in this decade.

Significant new research and development effort is necessary to overcome technological barriers to the creation of a generation of supercomputers that tests these technical limits. Computer manufacturers in the U. S. have neither the financial resources nor the commercial motivation in the present market to undertake the requisite exploratory research and development without partnership with government and universities.

Unless these barriers are overcome, the primacy of U. S. science, engineering, and technology could be threatened relative to that of other countries with national efforts in supercomputer access and development. Although the Federal Government is the first and by far the largest customer for supercomputers, there are no national plans to stimulate the development and use of advanced computer technology in the U. S.