# FCCSET Panel on Supercomputers 1982-1985

Formed to recommend what, if anything, the U. S. Government should do.
As a result of discussions, three interagency working Groups were formed:

- Procurement
- Access
- Research Coordination

The reports of these working groups led to the formation of FCCSET committees on high performance computing (incorporating the Access, Procurement, and Research groups).

P

January 1983

Report to

The Federal Coordinating Council

on Science, Engineering and Technology

Supercomputer Panel

on

Recommended Government Actions to

Retain U.S. Leadership in Supercomputers

## I. Executive Summary

### A. Background

In January 1983, a Federal Coordinating Committee on Science, Engineering and Technology (FCCSET) Panel on Supercomputers was formed to examine what, if anything, the U.S. Government should do to stimulate the use and development of supercomputers. As a result of the FCCSET panel discussions, three interagency working groups were formed. These groups are to examine and make recommendations on the following supercomputer issues:

1) What should the Government do to ensure that the U.S. retains its lead in supercomputers? This group is named the Procurement Group.

2) What should the Government do to make supercomputers available to more researchers, particularly in universities? This group is named the Access Group.

3) A third group was formed to provide coordination among Government funding agencies research contributing to the technology base.

This Report describes the findings of the Procurement Group. The Procurement Group consists of the following members:

James F. Decker, DOE (Chairman)
George Deskin, DCI
Leonard A. Harris, NASA
E. F. Infante, NSF
Robert E. Kahn, DOD
Joe Smagorinsky, DOC

This Group was assisted by a DOE/NASA working group with the following membership:

James Decker, (DOE)
Dave Nelson, (DOE)
Bob Greeves, (DOE)
Don Austin, (DOE)
Glenn Kuswa, (DOE)
Bill Bostwick, (DOE)
Kate Foley, (DOE)
Paul Schneck, (NASA)

Before developing their recommendations, these committees discussed many issues of supercomputer development, use, and markets with computer scientists from Government laboratories and universities. Discussions were also held with senior management of the supercomputer vendors.

Supercomputers are the most powerful general purpose computers available for large scale scientific computation. At the present time, the Cray-1, Cray-XMP, and the Cyber 205 are the most advanced machines available. Fifth Generation computers are computers based on artificial intelligence techniques and are not conventionally defined as supercomputers, even though they may require very fast processors and large memories. The development of Fifth Generation computers as defined by the Japanese was not considered by this group. Similarly, special purpose scientific computers, which perform very well on one or two specific problems, were not considered.

B. Findings

1. Supercomputers have become an essential scientific tool in many areas of Government research and development. The U.S. Government is the largest user of supercomputers, having purchased or leased over 50% of all the present class of supercomputers, Class VI machines, currently in service in the U.S. These machines are being used to model many complex physical phenomena found in nuclear weapons design, weather prediction, aerospace, magnetic and inertial confinement fusion, fundamental physics, and many other fields of research and development. Government use of supercomputers is expected to grow in the future as our ability to model complex phenomena improves. This improved modeling capability will result in continued replacement of expensive experimental testing.

2. The use of supercomputers in the private sector is increasing rapidly. Key industries such as oil, automobiles, electronics, chemicals, and aerospace are applying the supercomputer to technical problems that have important commercial implications. Future development of competitively priced commercial products with superior performance in these and many other industries will depend heavily on the availability of continually improved machines within a few years. The Japanese apparently have recognized this potential and have initiated a major supercomputer effort called the National Superspeed Computer Project. In addition to the potential international market, the Japanese are planning for a large domestic market where according to published reports, many of the supercomputers that they plan to produce will be sold for use in Japanese industry.

3. The U.S. Government and its laboratories historically have played an important role in causing each new supercomputer system to be developed. The Government has always been a "friendly" buyer by accepting certain risks associated with buying the first few machines of each new class. Each new machine has arrived at a Government laboratory with minimal operating software and devoid of any applications software. The laboratories, particularly Los Alamos and Livermore, have developed substantial amounts of software required to make these machines into useful systems.

4. Current Government programs require supercomputers with capabilities substantially in excess of those available today. Within this decade, programs in nuclear weapons design, aerospace, weather prediction, fusion research, and many areas of fundamental research and development require computers with capabilities at least 200 times greater than present Class VI machines, and requirements for even greater capability are projected for the next decade.

5. Because of the established relationship of supercomputers to leadership in both Government and private sector research and development, the U.S. can not afford to relinquish its lead in this important area of technology. Programs crucial for the defense and economic well-being of the country can not become dependent on foreign vendors and on the policies of their governments.

6. The U.S. supercomputer industry is small. Only two companies, Cray Research and Control Data Corporation, have produced the latest class of supercomputers. A third company, Denelcor, has entered the field. The industry is small because the development costs are large, the risks are high, and the market is currently limited to the most sophisticated applications. Each of the vendors are primarily manufacturers of computer mainframes, i.e., they produce little software and they do not manufacture components or peripherals (however, Cray has recently invested in their own integrated circuit facilities, and Control Data does make high performance disk drives). In the future, it is questionable whether the U.S. vendors can outperform foreign government supported national efforts by large integrated companies such as Hitachi, Fujitsu, and Nippon Electric Company, (see Appendix A). Nippon Electric has made the highest performance claims for a machine capable of 1,300 million floating point operations per second (MFLOPS), scheduled for delivery in 1985. The Cray-1 performs at peak speeds of about 100 MFLOPS and the Cyber 205 has possible peak speeds of 600-800 MFLOPS. The Cray-2, scheduled for delivery in early 1985, is projected to perform at a peak of 800 MFLOPS. Performance of these new machines on actual laboratory job mixes is, of course, an open question.

C. <u>Recommendations</u>

The U.S. should not relinquish its lead in this important area of technology. However, a "business as usual" approach on the part of both the vendors and the Government has a low probability of competing successfully with foreign joint government/industrial collaborations. During the past decade the capability of U.S. supercomputers has increased by a factor of only 4-5 every 4-5 years. The Japanese goal is a factor of 100 improvement by 1990. The U.S. Government should continue to play its traditional role in relation to the U.S. supercomputer vendors - i.e., the Government should continue to be a friendly buyer, develop software as necessary, and support basic computational sciences and base technology development. However, all activities associated with that role should be accelerated in a coherent manner in order to meet the challenge of foreign competition. Specific recommendations include the following:

1) The Government should set as a national goal the development of supercomputers with at least 200 times the capability of Class VI machines in this decade. The Government should provide an incentive by guaranteeing to buy at least three of each U.S. developed and manufactured supercomputer system that meets the goal. (Further work needs to be done to define more completely the 200x goal and the terms of a guaranteed buy). If successful, this process should be repeated for machines beyond the 200x class.

2) The Government should accelerate its purchases of supercomputers in order to ensure the health of the industry. Profits from sales of the present class of supercomputers will contribute to the development of the next class. The Government should continue to be a friendly buyer, particularly for prototype and production machines that represent important steps toward achieving the 200x goal.

3) It currently appears that the only feasible technical approach to meeting the 200x goal in the near term is to utilize parallel processors. Development of the software to obtain the theoretical benefits of using many processors in parallel to solve a single problem will be a significant challenge. Immediate experience with parallel processors is required. Government laboratories, and perhaps with government support a few universities, should purchase or lease several machines with interesting parallel architectures for experimentation. These "experimental" machines should be put on networks so that researchers from many institutions, including universities, can contribute to development of new languages, algorithms, software tools, and applications systems. A difficult but worthwhile goal is the development of a common operating system for the next class of U.S. supercomputers.

4) The development of high performance peripherals required for future generations of supercomputers is a well recognized problem. One of the primary difficulties is the small market that makes development of the desired hardware by normal private sector financing unattractive. Two possible solutions are a guaranteed government buy and/or government funding of research and development. A similar problem exists for high speed components. The vendors may not be able to obtain advanced high speed memory and logic circuits from U.S. suppliers in the future which could have serious implications for development of new machines. Continued government support of base IC technologies including promising new approaches to fast logic and memory circuits is important. The special requirements of the supercomputer industry should be examined in-depth.

Fields that have been and continue to be strongly dependent on computational modeling include nuclear weapons design, magnetic fusion energy, cryptographic analysis, aerodynamics, integrated circuit design, inertial confinement fusion, reactor safety, atmospheric research and weather forecasting, astrophysics, molecular biology and chemistry, and fundamental physics research.

The speed of computation that can be achieved in the fastest computers available at any given time has increased by some seven orders of magnitude from the early 1940s to the present. Three of these orders of magnitude resulted from the transition from electromechanical to electronic circuits. Another three orders of magnitude have been achieved in the speed of the electronic components (for example, the IBM 701 of the early 1950s had a cycle time of 12 microseconds, and the Cray-1 has a cycle time of 12 nanoseconds). The other order of magnitude in this overall increase has come about through increasing levels of parallel execution of computational tasks. Equally dramatic advances have come about through improvements to the algorithms used to solve complex problems.

The supercomputer industry is now facing a demand for several orders of magnitude increase in performance in this decade, but the projected performance increases available from faster components appear to be limited to at most one order of magnitude. Hence, it is clear that new and explicit forms of parallel computer organization will be required to meet the needs of the nation. Achieving success in this radical change in the design of supercomputers will require the cooperation of some of the best minds in the nation.

The early years of the history of supercomputers were characterized by a strong interaction among government, universities, and industry that guided and accelerated the development of very powerful scientific computers. The mathematician John von Neumann served as an important catalyst for this interaction: he was a consultant to both the government and the computer industry, and he led the development of the computer at the Institute for Advanced Study at Princeton that served as a model for the development of a whole generation of computers, including the IBM 700 series computers and their successors, the ILLIAC, the Johnniac, the MANIAC, the ORDVAC, and the ORACLE. This highly beneficial interaction among government, universities, and industry continued through the decades of the 1950s and 1960s, but suffered a serious decline in the 1970s. Universities in particular have become isolated from the mainstream of supercomputer research and development. The leading supercomputer of the early 1970s was the Control Data Corporation 7600, but none of these computers was installed at an American university (although several were installed at foreign universities). Even today, only three American universities have on-site supercomputers. This unfortunate situation has resulted in fewer graduates with backgrounds in large scale scientific computing and fewer research projects being conducted in universities in this field than in previous decades. There is general agreement among these communities that it is essential to reestablish their close collaboration of earlier decades for the United States to maintain its leadership in supercomputer development, marketing, and applications.

## III. Anticipated Requirements for Future Supercomputers

At the time this report was prepared, approximately 70 Class VI supercomputers have been sold in the U.S., the U.K., Germany, France, Sweden, and Japan. Only about 10% of these are at universities (three in the U.S.). Several reports have recently been written describing the need for more facilities in the academic environment (e.g., the Lax Report and the Press Report). Additional information on the applications in Government laboratories and industry is provided below.

### A. Applications in Government Research and Development

The Federal Government purchased 29 CDC 7600 supercomputers, about 40% of the total number built. Manufacture of the CDC 7600 ceased in 1977, yet all 29 remain in productive use. The current generation of supercomputers includes the Cray-1 and the CDC Cyber 205. The Government has already acquired 29 of these, and several more are expected to be acquired over the next two years, including the recently announced Cray-XMP.

The Department of Energy is the largest user of supercomputers, with 13 Cray-1's and two Cyber 205's already installed. Nuclear weapons design and engineering are heavily dependent on the use of supercomputers, as are magnetic and inertial confinement fusion research and nuclear reactor development. NASA has two Cray-1's and one Cyber 205; they are used for atmospheric modeling, aeronautical design, and scientific research. The National Science Foundation has two Cray-1's installed at the National Center for Atmospheric Research, where they are used for atmospheric and other scientific research. The Department of Commerce uses two Cyber-205's, and a third will soon be added for NOAA's weather forecasting and oceanographic research. The Department of Defense uses a Cray-1 for weapons effects simulation and a Cyber 205 for weather forecasting.

The graph (Fig. 1) shows the present and projected Government installed base of supercomputers. DOE's designation of classes has been used to label each succeeding class of supercomputers. Class IV installations after the introduction of Class VI machines are not included because, by definition, these are no longer considered supercomputers. (In fact, there are many more installed Class VI and Class V machines, such as IBM 3081's and Cyber 176's than are shown in the graph).

5) Government support of long range research and development for scientific supercomputing should be increased. This research is required to develop the new approaches in hardware and software necessary for the substantial increases in computing power required in the next decade. Since most of this research would be supported in universities, an important product of this research will be the trained supercomputer architects and developers of the future.

6) The Government should take action to greatly increase access to supercomputers for researchers, particularly in universities. Improved access will not only bring this modern scientific tool to universities and improve their competitive position in research, but it will also help train new scientists and engineers in the use of supercomputers.

7) Research and development tax incentives for the industry should be explored by a Government committee with members from OMB, Department of Commerce, and Department of Justice.

8) U.S. export control of supercomputers has a significant effect on U.S. supercomputer vendors since Western Europe represents 40% of the market. Acceleration of the export licensing process should be explored by DOD and Commerce along with other appropriate agencies so that U.S. vendors are not unduly penalized in competing for foreign customers. Of course, export control of supercomputers must remain consistent with national security considerations.

9) A permanent interagency group should be established to coordinate individual agency supercomputer activities as necessary to implement the above recommendations. This group should function at both the policy and technical level.

## II. Historical Perspectives

The federal government has exerted a powerful influence on the development of large scale scientific computers (supercomputers). The first electronic computer, the ENIAC, was built by the University of Pennsylvania under contract to the U.S. Army Ballistics Research Laboratory, and the first program run on the ENIAC was a nuclear weapons design calculation by the Los Alamos Scientific Laboratory. In fact, the Manhattan Project and the ENIAC project were initiated the same year (1943). In the succeeding four decades, the defense and computer communities have had a mutually beneficial relationship: defense programs have provided much of the motivation and funding for the development of supercomputers, and the computer industry has provided the tools for solving the highly complex problems that would otherwise have remained intractable.

Future projections are intentionally conservative on the graph. They are based largely on replacement of present machines following historical trend and new installations budgeted or planned. Not included is allowance for emerging applications and greater access, which could add another 10-20 machines by 1990. Planned new installations include NASA's Numerical Aerodynamic Simulation, DOE's Energy Research Network, NSF's supercomputer program, Commerce's National Bureau of Standards system, and DOD's Naval Research Laboratory.

Based on these projections, it appears that the direct Government requirement for supercomputers will be at least 30-40 machines from each new class. New applications could increase this to 50 Class VII computers.

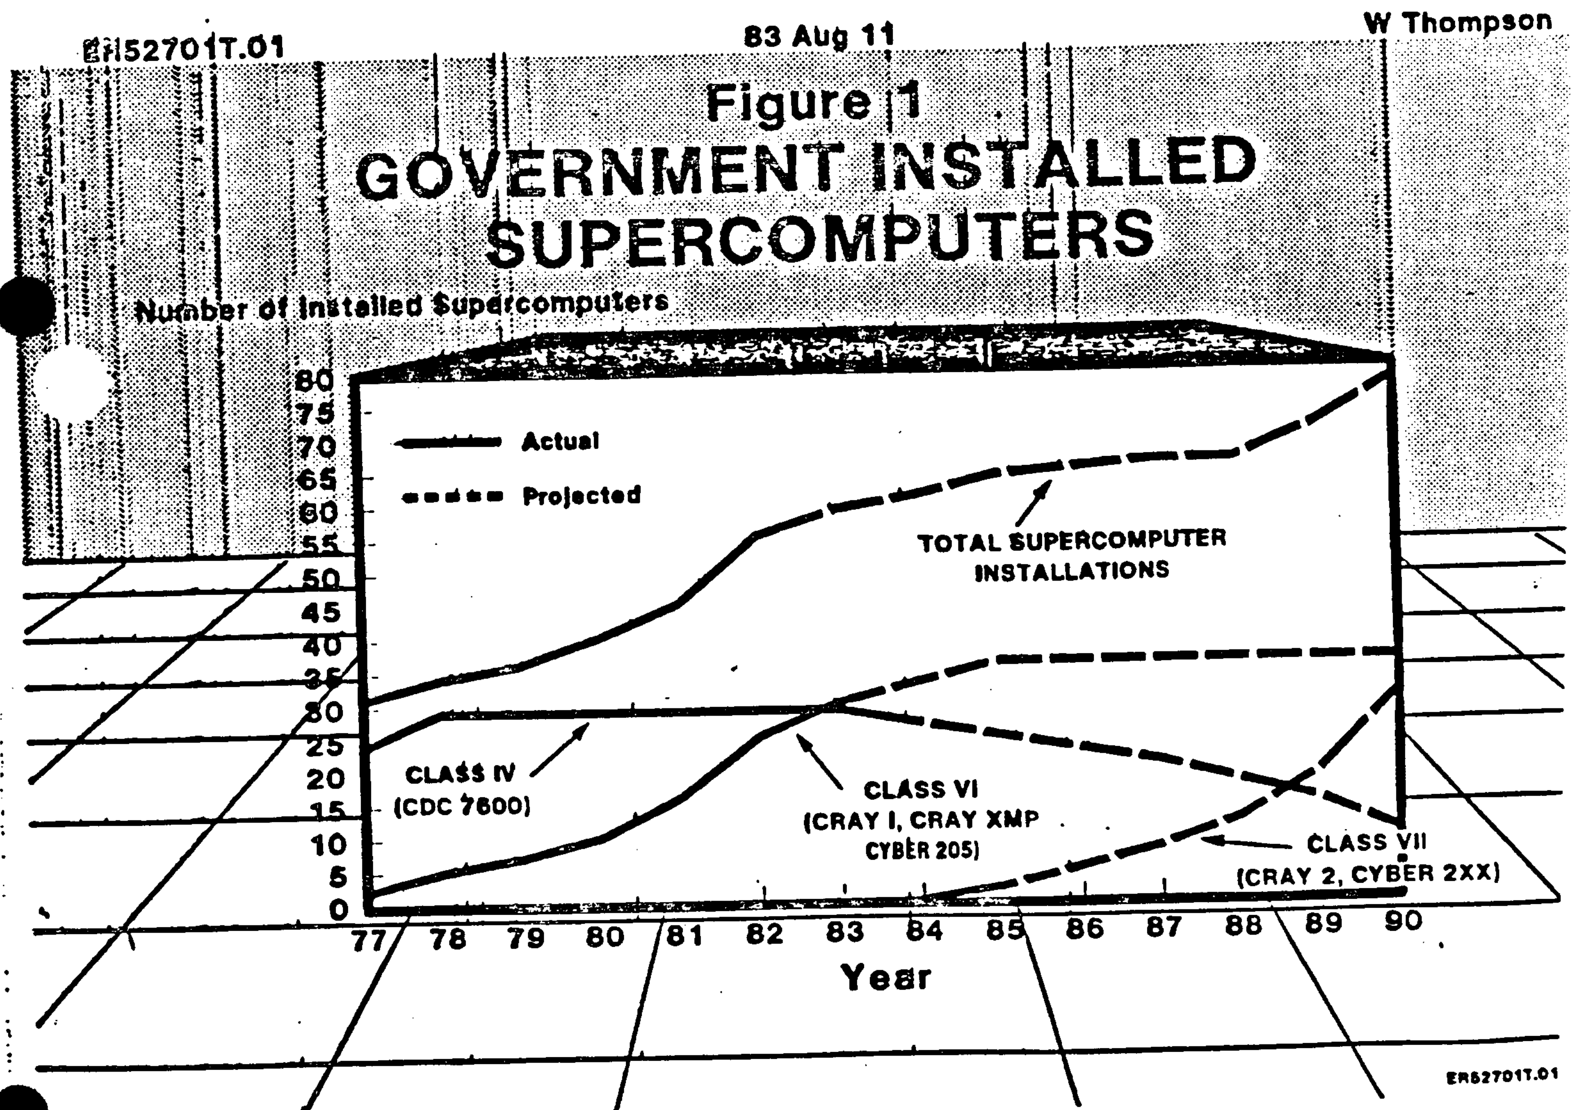

## B. Commercial Applications of Supercomputers

Scientific applications have been developed and perfected by the national laboratories and universities in conjunction with research and development activities. The applications are gradually adopted by industry when the techniques are demonstrated to be valuable in modeling complex processes. Programs are typically made available to industry at a nominal cost through distribution organizations, such as DOE's National Energy Software Center, NASA's Cosmic, and DOC's National Technical Information Service.

Examples of current applications in use by industry are:

Aerospace Industry: Simulation of air flow around aircraft using finite difference and vortex tracking methods; structural analysis using finite element methods.

Petroleum Industry: Seismic exploration using inverse scattering techniques; tertiary recovery of oil from reservoirs using front tracking methods.

Automotive Industry: Combustion modeling for internal combustion engines using flame front tracking and chemical kinetics programs; computer aided design and testing of automobile parts.

Electronics Industry: Computer aided design of VLSI circuits; numerical solution of large sets of differential equations describing circuits using numerical linear algebra techniques.

Engineering and Manufacturing: Computer aided design of machine parts; testing structural stress-strain properties of structure using finite element methods; non-destructive evaluation using image reconstruction techniques.

Power Industry: Modeling of nuclear reactor safety systems; power grid optimization using linear programming techniques.

Chemical Industry: Process control and optimization; simulation of polymers; reaction diffusion modeling and turbulence.

Movie Industry: Computer graphics techniques for image generation and animation.

## IV. Industry Outlook

A healthy U.S. supercomputer industry is necessary both for its effect on national defense issues and on the balance of trade. The computing equipment market, of which supercomputers are a bellwether, is one of the leading U.S. exports. Nuclear weapons design programs are highly dependent on advanced supercomputers and must utilize the highest capability available - this can only be assured if these are a U.S. product.

Supercomputers represent a small niche in the large computer market. The risks and costs associated with developing a new supercomputer system are large. Consequently, most large U.S. computer vendors have not entered the supercomputer field. An additional risk factor for the supercomputer vendors is their dependence for state of the art integrated circuits on merchant semiconductor manufacturers over whom they exert little, if any, influence. A similar situation exists for high speed peripherals.

In order to help remove some of the risk factor, there must be a well identified market - e.g., the Government sponsored market. If such a market is not clearly delineated, foreign manufacturers aided by their government's funding, and favorable export policies, will be the only companies willing to risk precious capital resources on poorly defined markets.

## V. Goals

Because of the strategic importance of supercomputers to the long term strength of the United States, both economically and militarily, it is imperative that an adequate supply of domestically produced supercomputers be available in the U.S. Toward this end, the goal of U.S. Government policy should be:

<u>To retain U.S. leadership in the development, production, and application of supercomputers</u>.

The implementation of a policy based on this goal will be instrumental in:

1. Meeting U.S. government needs for advanced supercomputers and for the trained scientific and engineering personnel necessary for their use;

2. Meeting U.S. industry needs for advanced supercomputers for oil exploration, aircraft design, integrated circuit design, automobile engine design, etc.;

3. Maintaining this country's strength and vitality in supercomputer technology through support of basic research and education;

4. Providing incentives necessary to accelerate the development of advanced supercomputers by U.S. manufacturers; and

5. Retaining the U.S. lead in computer related technology as a whole through spin off from leading edge supercomputer technology.

## VI. Recommendations

The Government has traditionally been a major purchaser of supercomputers and has supported hardware and software R&D that benefits supercomputer development. The intent of these recommendations are: (1) to focus Government activities sufficiently to provide incentives for accelerating supercomputer development, and (2) to focus Government-sponsored R&D in generic areas that are crucial for improved supercomputers but inadequately stimulated by market forces in the private sector. The development and production of complete supercomputer systems, including architecture, hardware, and software, cannot now be accomplished by any one organization. The best talents of U.S. industry, Government laboratories, and universities must be brought to bear in order to ensure expeditious development of complete, usable systems.

The recommendations are designed to address both the short-term and long-term requirements for accelerating supercomputer development as necessary to provide the large scale scientific computing needs of Government and to maintain U.S. leadership in this industry. The recommendations deal with several aspects, including manpower training, expansion of supercomputing applications, development of several 200x supercomputers by 1990, and long range R&D required for computers with capabilities beyond 200x. A 200x supercomputer refers to a system having a capability that is about 200 times that of current Class VI systems.

GOVERNMENT ROLE: Most of the recommendations provide recognition of the role that the government has played historically in supercomputer development rather than suggesting completely new government approaches. However, the recommendations do call for an increased emphasis on and an acceleration of this traditional government role in order to meet the future needs of the country and to retain U.S. leadership in this field.

1. Recommendation: Stimulate industrial development of several 200x scientific supercomputer systems by 1990.

The development of 200x supercomputers by 1990 is a challenging yet technically feasible goal. Several manufacturers have defined architectures that can be extended to this level of performance, but new development will be required in several areas: fast circuits with large scale integration, high density component packaging with adequate cooling, high performance peripherals, system integration of multiple processors, and improved software. Traditionally, supercomputer vendors have developed new architectures and have packaged components. Software development has been assisted by Government laboratories, with the vendors supplying only rudimentary software. Integrated circuits using off the shelf technology have been purchased from IC vendors. High performance peripherals have been developed and manufactured by vendors. The goal of the recommended program is to use the best capabilities of the supercomputer vendors, the IC vendors, peripheral vendors, Government laboratories, and universities.

The heart of this strategy is the guaranteed Government purchase or lease of three of the first supercomputers from each vendor who meets the 200x goal in the FY 1986-1990 time frame. The intent of this guarantee is to provide sufficient incentive for vendors to accelerate supercomputer development beyond current trends. The objective is to produce computer systems that not only meet Government requirements, but can also compete in the commercial market at a system price of $10-20 million.

The 200x goal is not well specified at this time and needs to be examined by an appropriate group of technical experts. Traditionally, new machines are "benchmarked" on a package of Fortran routines to be representative of the typical workload of existing large scale scientific computer centers. A recent suggestion is to judge supercomputer system performance on solving four or five scientific problems of national importance rather than using benchmark codes. The goal using this approach would be to execute the problem benchmark set at a speed-up of some 200 times that currently available, allowing for improvement in component speed, concurrent processing architectures and more efficient algorithms and languages. This new approach should be seriously considered.

In addition to the guaranteed acquisitions there are several other considerations in the recommended strategy.

a. Architecture and Packaging

The supercomputer vendors have the expertise to develop new architectures and to package components for the 200x systems, although this goal will provide challenges to the vendors in these areas. The guaranteed acquisition plus the remainder of the supercomputer market is considered to be an adequate incentive. No direct Government support for vendor R&D in these areas is recommended.

b. Integrated Circuit Development

Acquiring fast logic and memory circuits for advanced supercomputers has become a problem for U.S. vendors. IC manufacturers in the U.S. are not developing on a timely basis the specific technology required for supercomputers because this market is very small compared with the market for high volume products such as personal and business computers. Japanese IC manufacturers have put more effort into supercomputer technology; as a result, U.S. supercomputer vendors are already forced to obtain some of their memory IC's from Japan. This leaves the U.S. vulnerable to the trade policies of Japan, or of individual IC manufacturers who may refuse to export new IC technology until internal needs are satisfied. Importing logic chips is particularly dangerous, because the internal logic design of U.S. computers must be revealed in order to produce these chips.

Because IC development is so crucial to advancing supercomputer performance, and reliance on foreign vendors is dangerous, U.S. capability in this area must be examined carefully. At present, no U.S. supercomputer vendor is able to produce its own IC's. Cray Research has recently invested in setting up its own chip manufacturing facilities. However, high speed IC manufacturing is a most difficult undertaking, and many companies have failed in similar attempts. For these reasons, Government should support the base technologies associated with high speed IC development. R&D should be supported on silicon devices, advanced semiconductors such as gallium arsenide and new promising technologies. Cost to the Government could be kept to a minimum by piggybacking on existing Government sponsored or private activities. The component question should be examined in more depth.

c. High Performance Peripherals Development

Peripherals development suffers from the same problems as IC development. The market for high performance peripherals, such as disk storage, is relatively small. Consequently, they attract little development funding within the manufacturers. Recent experience in high performance peripherals development indicates that common supercomputer peripherals could be specified by the industry, resulting in a truly generic development program. Since advanced peripherals are essential to future supercomputer systems and market forces do not adequately push their development, some Government support for their timely development should be strongly considered. Two possible mechanisms are guaranteed government buys and government funding of research and development.

d. Software

Software requirements for future supercomputers include (1) improved operating systems and utilities to effectively use parallel processing, (2) an improved Fortran-like language with parallel capability, and (3) new algorithms to apply parallel processing techniques to single, large problems. Supercomputer vendors currently supply only rudimentary software; they receive considerable software assistance from Government laboratories. It is unlikely that vendors will be able to develop effective software for future supercomputers without continued assistance, yet software quality will be an important factor in international competition.

The strongest existing groups for developing supercomputer operating systems are in the national laboratories. Their expertise should be fully utilized. Alternatives should be explored to best use the resources of vendors, laboratories and unviersities to develop and maintain superior software.

A single common operating environment for supercomputers, including the operating system, utilities, and a Fortran-like language, would reduce the resources required by industry to develop software and would enhance competition by simplifying installation of supercomputers from multiple vendors at a single site. Discussions with supercomputer vendors indicate that, in principle, they welcome the idea of a common operating environment. The government should take a more assertive role leading to a common operating system and a high level user language.

Research on algorithms is required to enhance our ability to use parallel processing techniques in applications software. This research would benefit from the availability of experimental and prototype computers to researchers at government laboratories and at universities. Support of research of this type by the government should be accelerated. In the long term, vendors should be encouraged to produce high quality software.

e. Prototype Acquisition

Several vendors have announced supercomputers whose architecture may be extendible to 200x capability. These computers can legitimately be considered incremental steps leading to the 200x machine. Early procurement by the government should be assured to speed software development and provide feedback to the vendor as the 200x machines are being developed. The Government has traditionally been a "friendly buyer" of the first copies of supercomputers. This role is now even more important to encourage accelerated supercomputer development.

Purchase of prototype computers from domestic producers as soon as they become available will make increased capability available earlier and provide closer ties to the needs of mission oriented Government agency users. Their purchase will also provide incentive by providing partial early return on the investment for the development of hardware and software systems. In addition to purchases of prototypes, Government purchase of production versions of machines providing significant incremental advances will help maintain a financially healthy industry that can afford to invest in development of the 200x machine and even more advanced systems.

This type of advanced procurement notice will encourage industry research and risk-taking necessary to develop breakthrough units to achieve large increment increases in performance. Mission oriented agencies must work with computer developers to focus designs to meet critical government requirements in scientific, space, and military applications.

GOVERNMENT ROLES: (1) Provide support for supercomputer development through multiple (at least three) procurements of the first 200x units and through support of essential R&D in components, peripherals, and software; and (2) Provide support for supercomputer development through purchase of new prototype and production supercomputers which are significant steps toward the 200x goal.

2. Recommendation: Accelerate hardware and software development through increased Government support of experimental supercomputer systems.

The purpose of the program in experimental computers is to provide an early opportunity for government laboratories and universities to explore new architectures and advanced concepts for a range of applications or to accomplish a unique task specific to the user organization. Experimental computers are defined as systems with hardware that is thoroughly debugged and warranted by the vendor, but with operating and support software that is not mature and in some cases may be totally lacking. Such systems are procured on a research and development basis rather than in response to program workload requirements, and will be used to support laboratory and university research on architecture, algorithms, and develoment of software rather than to solve production workload problems. Experimental computers may be a precursor to future production systems, but may differ in terms of engineering and performance from later production versions.

Immediate experience with experimental processors is required to determine how to use new architectures and to guide specification and evaluation of new systems such as the 200x supercomputers. Several machines have been developed recently which, although not supercomputers, have very interesting architectures and present opportunities to learn about parallelism concepts in particular. Several of these systems should be installed in government laboratories and universities and made available, through computer networks to researchers for exploration of new algorithms, languanges, software tools, and applications.

Care should be taken to assure that a clear distinction is made between the acquisition and use of research and experimental systems versus the acquisition and use of prototype and production computers, in order to avoid misunderstandings concerning the expectation of results.

GOVERNMENT ROLE: Provide increased resources for and access to advanced research and experimental supercomputer systems.

Increased funding is required for experimental system developments and for networks for access to these systems by the Government, university, and computer industry R&D community.

3. Recommendation: Provide improved access for expansion of applications and manpower training.

At the present time, very few universities have access to supercomputers. Because of this lack of access, there is very little training of university students in the use of supercomputers. Consequently, government laboratories and industry have to expand substantial time, one to two years, training new hired scientists and engineers to use large scientific computers. Further, the lack of access of university researchers to supercomputers means that a large number of scientists do not utilize one of the most powerful scientific tools available today. In contrast, some European countries have placed new supercomputers in their universities, and U.S. researchers are complaining about their inability to compete in forefront research with their better equipped European colleagues.

Providing universities access to supercomputers as soon as possible is of paramount importance to the country's R&D activities as well as to manpower training. In addition, the availability of supercomputers to the many bright, young students and faculty would undoubtedly result in unanticipated applications for scientific computing. These new applications could broaden the market, increase supercomputer sales in the future, and strengthen the ability of U.S. industry to compete in worldwide markets.

GOVERNMENT ROLE: The mechanism by which access should be provided is the subject of an interagency committee and will not be addressed in detail here. However, the access question is considered to be of the utmost importance to the health of the supercomputer industry.

4. Recommendation: Enhance long range basic R&D for scientific supercomputing.

Supercomputer architecture ideas for the 1990's must come from research and development activities of the 1980's. Currently there are several interesting architectural studies being pursued in universities and in Government and industry laboratories. The areas under investigation include problem decomposition techniques, algorithms, languages, and software environments for massively parallel computer architectures. These new architectures, employing hundreds or thousands of fast processors working concurrently to solve a single problem, require a thorough reconsideration of methods for the decomposition of problems into parallel operations, the creation of numerically stable algorithms for solving the problem components, the development of new software tools for ensuring the robustness and correctness of the implementation, and the development of new techniques for representing the solutions.

The R&D projects envisioned include interdisciplinary teams of computational scientists and computer scientists working on all aspects of large scale scientific computing problems. The design and construction of several potentially strong candidate experimental machines will be an important proof of concept activity.

University researchers will play the major role in generating ideas and experimental software and in training graduate students in generating new architectural concepts and applications. Government laboratory staff are in the forefront of tackling real world, large scale scientific problems and have unique resources for participating in these research projects. Industry likewise has a unique role in providing state of the art production and testing facilities and would stand to reap great benefits in understanding future architecture and software issues that tend to limit industry use of supercomputers currently. The transfer of technology from the academic and laboratory research environment to industry would be as rapid as possible through these cooperative projects.

GOVERNMENT ROLE: Government agencies responsible for supporting basic research should expand significantly their support of computational aspects of research programs in the physical, mathematical and social sciences. In particular, the present research programs in computational mathematics, algorithms, software engineering, and the development of advanced experimental machines based on novel architectural ideas should be enhanced to provide the future advances in computational science and engineering. Both the research aspects and the training of graduate students and postdoctoral researchers should be emphasized.

This program will be coordinated by the participating agencies to share common facilities where possible, such as VLSI design and fabrication facilities, experimental machines, and computer networks.

5. Recommendation: Research and development tax incentives for industry involved in the development of supercomputer systems components should be explored by an appropriate group, e.g., a committee with representatives from OMB, the Department of Commerce and the Department of Justice.

6. Recommendation: Some process whereby the acceleration of export licensing can be achieved should be explored by an appropriate group so that U.S. vendors are not unduly penalized in competing for foreign customers.

7. Recommendation: A permanent interagency group should be established to coordinate individual agency activities in supercomputing as necessary to implement the recommendations in this report.

All Federal agencies directly involved in supercomputing R&D should be represented in the coordinating group, which will report directly to FCCSET. This group should function at both the policy and the technical level, and act to coordinate procurement of new experimental supercomputers, software development, networking and access, and long range R&D.

FCCSET Working Group

Dr. James F. Decker, Chairman
Department of Energy

Dr. Robert E. Kahn
Department of Defense

Dr. E. F. Infante
National Science Foundation

Dr. Leonard A. Harris
National Aeronautics & Space Administration

Dr. Joe Smagorinsky
Department of Commerce

Mr. George Deskin
Director, Planning & Policy Staff
Intelligence Community Staff

## Appendix A

The Japanese have begun a major effort to become the world leader in supercomputer technology, marketing, and applications. This effort includes not only the development of supercomputers by the three largest Japanese computer firms but the support by the Ministry of International Trade and Industry (MITI) of three national projects affecting supercomputer technology.

### 1. NATIONAL PROJECTS

#### 1.1 The Super-Speed Computer System (SSCS)

Six major Japanese computer firms (Fujitsu, Hitachi, NEC, Mitsubishi, Oki Electric, and Toshiba) are collaborating under the leadership of the Electro-technical Laboratory (ETL) to develop the technology that will make it possible to build a supercomputer capable of performing at a peak rate of 10 billion floating-point operations per second (FLOPS). This project is supported by about $100 million from MITI and is to be completed in 1989. If successful, this project would lead to the development of computers some 100 times as powerful as the Cray-1. The six companies involved in this project have revenues that are about 8 times as large as the revenues of the American supercomputer vendors; with the collaboration of a national laboratory (ETL) and the financial backing of MITI, the resources backing this project can be estimated to be at least 10 times as great as those available to the American supercomputer vendors.

#### 1.2 The Fifth-Generation Computer System (FGCS)

Eight organizations (Fujitsu, Hitachi, NEC, Mitsubishi, Oki Electric, Toshiba, ETL, and Nippon Telephone and Telegraph (NTT)) are collaborating in the research and development of technology intended to revolutionize the design of computers using artificial-intelligence concepts. A new institute, the Institute for New-General Computer Technology (ICOT) has been chartered to oversee this project. Total funding for the project has not been specified, but the funding for the first five years of this ten-year project will be about $430 million.

#### 1.3 New-generation technologies project

This project includes much more than is of interest for supercomputers, but it includes the development of high-speed components, three-dimensional device geometries, and components able to withstand hostile environments. The high-speed components for the supercomputer project include not only extensions of silicon technology but efforts to develop high-speed gallium arsenide (GaAs), high-electron-mobility transistors (HEMT), and Josephson Junction technology.

3. Recommendation: Provide improved access for expansion of applications and manpower training.

At the present time, very few universities have access to supercomputers. Because of this lack of access, there is very little training of university students in the use of supercomputers. Consequently, government laboratories and industry have to expand substantial time, one to two years, training new hired scientists and engineers to use large scientific computers. Further, the lack of access of university researchers to supercomputers means that a large number of scientists do not utilize one of the most powerful scientific tools available today. In contrast, some European countries have placed new supercomputers in their universities, and U.S. researchers are complaining about their inability to compete in forefront research with their better equipped European colleagues.

Providing universities access to supercomputers as soon as possible is of paramount importance to the country's R&D activities as well as to manpower training. In addition, the availability of supercomputers to the many bright, young students and faculty would undoubtedly result in unanticipated applications for scientific computing. These new applications could broaden the market, increase supercomputer sales in the future, and strengthen the ability of U.S. industry to compete in worldwide markets.

GOVERNMENT ROLE: The mechanism by which access should be provided is the subject of an interagency committee and will not be addressed in detail here. However, the access question is considered to be of the utmost importance to the health of the supercomputer industry.

4. Recommendation: Enhance long range basic R&D for scientific supercomputing.

Supercomputer architecture ideas for the 1990's must come from research and development activities of the 1980's. Currently there are several interesting architectural studies being pursued in universities and in Government and industry laboratories. The areas under investigation include problem decomposition techniques, algorithms, languages, and software environments for massively parallel computer architectures. These new architectures, employing hundreds or thousands of fast processors working concurrently to solve a single problem, require a thorough reconsideration of methods for the decomposition of problems into parallel operations, the creation of numerically stable algorithms for solving the problem components, the development of new software tools for ensuring the robustness and correctness of the implementation, and the development of new techniques for representing the solutions.

The R&D projects envisioned include interdisciplinary teams of computational scientists and computer scientists working on all aspects of large scale scientific computing problems. The design and construction of several potentially strong candidate experimental machines will be an important proof of concept activity.

University researchers will play the major role in generating ideas and experimental software and in training graduate students in generating new architectural concepts and applications. Government laboratory staff are in the forefront of tackling real world, large scale scientific problems and have unique resources for participating in these research projects. Industry likewise has a unique role in providing state of the art production and testing facilities and would stand to reap great benefits in understanding future architecture and software issues that tend to limit industry use of supercomputers currently. The transfer of technology from the academic and laboratory research environment to industry would be as rapid as possible through these cooperative projects.

GOVERNMENT ROLE: Government agencies responsible for supporting basic research should expand significantly their support of computational aspects of research programs in the physical, mathematical and social sciences. In particular, the present research programs in computational mathematics, algorithms, software engineering, and the development of advanced experimental machines based on novel architectural ideas should be enhanced to provide the future advances in computational science and engineering. Both the research aspects and the training of graduate students and postdoctoral researchers should be emphasized.

This program will be coordinated by the participating agencies to share common facilities where possible, such as VLSI design and fabrication facilities, experimental machines, and computer networks.

5. Recommendation: Research and development tax incentives for industry involved in the development of supercomputer systems components should be explored by an appropriate group, e.g., a committee with representatives from OMB, the Department of Commerce and the Department of Justice.

6. Recommendation: Some process whereby the acceleration of export licensing can be achieved should be explored by an appropriate group so that U.S. vendors are not unduly penalized in competing for foreign customers.

7. Recommendation: A permanent interagency group should be established to coordinate individual agency activities in supercomputing as necessary to implement the recommendations in this report.

All Federal agencies directly involved in supercomputing R&D should be represented in the coordinating group, which will report directly to FCCSET. This group should function at both the policy and the technical level, and act to coordinate procurement of new experimental supercomputers, software development, networking and access, and long range R&D.

FCCSET Access Group

James F. Decker
Department of Energy

Robert E. Kahn
Department of Defense

Larry Lee
National Science Foundation

Richard Dunlavey
NBS/ICST
Department of Defense

Ken Wallgren
National Aeronautics and Space Administration

George Rogers
Intelligence Community Staff

## Appendix A

The Japanese have begun a major effort to become the world leader in supercomputer technology, marketing, and applications. This effort includes not only the development of supercomputers by the three largest Japanese computer firms but the support by the Ministry of International Trade and Industry (MITI) of three national projects affecting supercomputer technology.

### 1. NATIONAL PROJECTS

#### 1.1 The Super-Speed Computer System (SSCS)

Six major Japanese computer firms (Fujitsu, Hitachi, NEC, Mitsubishi, Oki Electric, and Toshiba) are collaborating under the leadership of the Electro-technical Laboratory (ETL) to develop the technology that will make it possible to build a supercomputer capable of performing at a peak rate of 10 billion floating-point operations per second (FLOPS). This project is supported by about $100 million from MITI and is to be completed in 1989. If successful, this project would lead to the development of computers some 100 times as powerful as the Cray-1. The six companies involved in this project have revenues that are about 8 times as large as the revenues of the American supercomputer vendors; with the collaboration of a national laboratory (ETL) and the financial backing of MITI, the resources backing this project can be estimated to be at least 10 times as great as those available to the American supercomputer vendors.

#### 1.2 The Fifth-Generation Computer System (FGCS)

Eight organizations (Fujitsu, Hitachi, NEC, Mitsubishi, Oki Electric, Toshiba, ETL, and Nippon Telephone and Telegraph (NTT)) are collaborating in the research and development of technology intended to revolutionize the design of computers using artificial-intelligence concepts. A new institute, the Institute for New-General Computer Technology (ICOT) has been chartered to oversee this project. Total funding for the project has not been specified, but the funding for the first five years of this ten-year project will be about $430 million.

#### 1.3 New-generation technologies project

This project includes much more than is of interest for supercomputers, but it includes the development of high-speed components, three-dimensional device geometries, and components able to withstand hostile environments. The high-speed components for the supercomputer project include not only extensions of silicon technology but efforts to develop high-speed gallium arsenide (GaAs), high-electron-mobility transistors (HEMT), and Josephson Junction technology.

## 2. COMMERCIAL PROJECTS

### 2.1 Fujitsu

Fujitsu has announced the VP-200 and VP-100 computers that will be available for delivery in the fourth quarter of 1983 in Japan; marketing of these computers will occur outside of Japan in 1984. These computers have peak-performance rates of about 500 and 250 million FLOPS, respectively; this compares to about 400 million FLOPS for American supercomputers. These computers will execute the IBM instruction set.

### 2.2 Hitachi

Hitachi has announced the S810-20 and S810-10 computers that will be available for delivery in the fourth quarter of 1983 in Japan and in other nations in 1984. These computers have peak-performance rates of about 600 and 300 million FLOPS, respectively, and will execute the IBM instruction set.

### 2.3 Nippon Electric (NEC)

NEC has announced that they will deliver two supercomputers in the first quarter of 1985: the SX-2, and the SX-1, with peak performances of 1300 and 650 million FLOPS, respectively. These computers are not IBM-compatible.

## 3. SITUATION AUDIT

### 3.1 Japanese strengths

In-house semiconductor capability. Hitachi, NEC, and Fujitsu rank 3rd, 4th, and 8th in the volume of semiconductors manufactured in the world. This gives these companies an in-house components development capability that can respond quickly to the needs of their systems designers for high-speed logic and memory components, and it will allow these companies to shorten the design cycle for new generations of supercomputers compared to their American competitors. This is a capability lacked by the American supercomputer vendors and is a key reason why American supercomputers have evolved so slowly in the last decade. The lack of an in-house high-speed components capability is the Achilles heel of the American supercomputer effort. Cray research is attempting to develop such a capability, but to expect this small company (1/10th the size of Hitachi, for example) to compete with some of the world's leading merchant semiconductor houses is unreasonable.

IBM-compatibility. Two of the three announced Japanese supercomputers are capable of running IBM-compatible software. This is an option available to supercomputer customers that has hitherto been lacking. This will meet the needs of organizations where the computer centers use IBM-compatible systems and also need supercomputer speeds in some of their programs. There is an estimated $200 billion worth of IBM-compatible software available in the world and this will make a software-rich supercomputer environment available to supercomputer users for the first time. It is entirely possible that a significant fraction of the forecasted growth in supercomputer usage could be in the IBM-compatible area. Because no American company offers an IBM-compatible supercomputer, this market will go by default to the Japanese unless some action is taken by American vendors to correct this deficiency in their project offerings.

Nationally subsidized supercomputer R&D. The risks associated with both the Super-Speed Computer System and the Fifth-Generation Computer System are being shared by the Government of Japan and the Japanese computer industry in the form of Government subsidies. In the long term, it is unlikely that the relatively small American supercomputer industry can compete successfully against an effort having the backing of the Government of Japan and an order of magnitude more resources than are available to the American supercomputer vendors.

### 3.2 Japanese weaknesses

Lack of experience with supercomputers. The Japanese have two American supercomputers installed and two of the supercomputer companies have previously developed vendor processors, so they are not without experience in the field. However, compared to the long history of design, development, and application of supercomputers in the U.S., the Japanese are initially at a decided disadvantage. However, the ability of the Japanese to do "reverse engineering" on competing products has been demonstrated many times, and the Japanese relative lack of experience will be a short-lived advantage for the Americans.

Lack of an innovative tradition. The Japanese have moved quickly into the supercomputer market by following the lead of the American designers and attempting to develop a superior product on the basis of the American designs. However, the next generation of supercomputers will require true innovation in the development of parallel processors, for which the Japanese cannot simply copy American designs. It remains to be seen how well the Japanese meet this challenge.

Weak higher education system. Compared to the American and European systems of higher education, the Japanese universities are decidedly weak in computer science education. The Japanese universities graduate more engineers than do the Americans, but the quality of computer-science education is higher in the United States than in Japan. This is an advantage that clearly must be exploited in American planning.

Lack of a supercomputer marketing base. The Japanese do not have a supercomputer marketing base, i.e., an infrastructure of sales, consulting, and maintenance personnel. A supercomputer by itself is of no use: it must have extensive support, and the Japanese have yet to develop that support system.

3.3 Japanese opportunities

Minimal competition. The supercomputer industry differs from most areas of the computer industry in that there is a small market and only minimal competiton for that market. This makes the American supercomputer vendors vulnerable to a concerted effort by a foreign vendor to marshall larger resources than are available to the American vendors and overwhelm them in the long term.

IBM-compatible supercomputer market. The last IBM-compatible supercomputer-class machine was the IBM 370/195, marketed in the early 1970s. Since then, the customer who needed supercomputer speed bought that speed at the expense of a software-poor computing environment, because the market forces were not strong enough to pay for the development of extensive supercomputer software. IBM has ignored this market, apparently for two reasons: (1) it was only a small market, and (2) it was contrary to IBM's and the industry's interest for the supercomputer vendors to be overwhelmed by the IBM giant. Further, IBM customers, because they were willing to choose the high-end IBM machines (such as the 308X) in order to retain the advantages of IBM-compatible software. With the entry of the Japanese supercomputer offering, the customer can now have both supercomputer speeds and software richness.

Internal Japanese requirements. A part of the Japanese motivation for developing their own supercomputers is to meet the needs of their own users. If ten Japanese customers bought American supercomputers, then some $100 million would flow out of Japan. Spending $100 million to develop their own competing products would save that investment for Japanese companies. Further, Japan is not just interested in the supercomputer market as such; rather they are interested in supercomputer applications, including aerospace, energy systems design, weather forecasting, integrated circuit design, fluid dynamics, structural design, reactor safety, and exploration for natural resources.

4. RESPONSES

In addition to the responses generated in the U.S. by these Japanese actions, there have been responses in Western Europe.

## 4.1 European Economic Community (EEC)

The EEC is completing plans to begin a program called "ESPRIT" (European Strategic Plan for Research in Information Technology) that calls for collaboration in "pre-competitive" R&D, very much in the Japanese style. This project is to run five years and be funded at about $1.5 billion, with about half of the funding to come from the governments in the EEC and half from the participating companies. This project is to begin in January 1984.

## 4.2 Great Britain

Great Britain has undertaken a project in response to the Japanese fifth-generation project, to be funded at about $550 million over five years, with the funding costs to be shared by the British Government and private companies. This project will focus on the following: (a) software engineering, (b) man-machine interface, (c) artificial intelligence, and (d) VLSI. Britain will also participate in the ESPRIT project.

## 4.3 France

France has passed a "Science Program Law" that calls for increasing the level of national support from 1.8% of GNP to 2.5% of GNP by 1985 (this is about the level of expenditures in the U.S. and other western nations, and Japan). This funding will be used to increase funding and staffing at national research centers and for grants, loans, and tax exemptions for private industry. France is considering building their own supercomputers, but no decision on this has been announced yet. France will participate in the ESPRIT program.

## 4.4 Germany

In addition to participating in the ESPRIT program, Germany has recently increased its support for parallel processing by $4 million annually.

## 4.5 Other

Both Korea and Taiwan have announced plans for government support of collaborative R&D programs.

# 5. SUMMARY

The scope of the Japanese supercomputer effort is much larger than the scope of the American supercomputer effort. This larger set of resources, if applied over the long term, will simply overwhelm the relatively small American supercomputer effort that currently exists unless there is some appropriate response from the American supercomputer industry, the government, and academia.

# FY 1985 Annual Report of the Federal Coordinating Council on Science, Engineering and Technology (FCCSET) by the FCCSET Committee on High Performance Computing: Procurement/Access Group (Committee)


## Summary

During the past year, the Committee met on a regular basis to review government and industry supported programs in research, development, and implementation of new supercomputer technology. The Committee maintains an overview of commercial developments in the U.S. and abroad. It regularly receives briefings from Government agencies to facilitate interagency communication on Government agency sponsored R&D efforts and makes such information available, where feasible, to industry and universities. In addition, the committee coordinates agency supercomputer access programs and promotes cooperation with particular emphasis on aiding the establishment of new centers and new communications networks.

Training of both operators of new supercomputer centers and new users has been a major concern. As a result of Committee discussions, training programs for new users were initiated this summer by the National Science Foundation and were funded by NSF and DOD. Through the Committee's efforts, existing supercomputer centers such as those operated by DOE, NASA, and NSF (NCAR), were encouraged to help train the operators of newly established centers. Such help is being provided.

An interagency MOU has been drafted for the purpose of trading small amounts of supercomputer resources among agencies. The primary reason for trading computer time is to promote efficient use of Federal resources and to increase the productivity of users.

The Committee made its annual visit to vendors in May and found that substantial progress had been made by Cray Research and ETA Systems toward developing their next generations of machines. Technical progress had also been made by Denelcor; however, the company was facing severe financial difficulties. The supercomputer vendors are still having difficulty in obtaining high performance IC's from U.S. chip makers leaving them dependent on Japanese suppliers. In some cases, the Japanese chip suppliers are the same companies, e.g., Fujitsu, that provide the strongest foreign competition in the supercomputer market. There is some evidence that the Japanese companies are not shipping their latest state-of-the-art components to their U.S. customers. This represents a serious problem if U.S. supercomputer vendors are forced to use IC's with performance inferior to their main competitors. The Committee believes that this is a particularly unhealthy situation and will continue to examine the issue with the intention of making recommendations for solving this problem.

Another continuing technology problem is the development of high performance peripherals required to match the capabilities of new supercomputer systems. Of particular concern is disk storage. Again the primary problem is the small market for supercomputer peripherals. Some progress has been made in establishing communication between supercomputer vendors and the Government sponsored work on magneto-optical storage being conducted by RCA and 3M. This development promises the large storage capacities and transfer rates required for supercomputer systems.

During the past few months, Denelcor's financial condition deteriorated further and the company is now essentially out of business. Denelcor had produced the HEP-1 computer based on an innovative parallel architecture, but its relatively slow technology limited its performance to below the supercomputer range. Their plans for the HEP-2 and HEP-3 might have resulted in supercomputers with different architectures than the planned 4-16 vector processor Cray and ETA machines. There are classes of problems that probably would have run very efficiently on the Denelcor machines.

Annual Visit to Industry

The Committee annually visits supercomputer manufacturers in order to maintain the established dialog, to observe R&D activity, to obtain information on new systems and to gain a better understanding of industry problems. Since the Committee's last report to you, the Committee visited Cray Research, Inc., ETA Systems, Inc., and Denelcor, Inc. in 1984 and 1985. At each company, proprietary performance forecasts or goals of new systems and products were presented. A typical presentation by manufacturers included the following:

> Progress in developing their next supercomputer and any plans for machines beyond the next one that they wish to share; any bottlenecks to advanced machine development that are of particular concern to them; their views of the Japanese supercomputer efforts; recommendations for Government actions related to supercomputers; topics or problems they wish to raise.

Such meetings give the Committee perspective about future technology, possible issues for the Government to address and an up-to-date perception of industry thinking.

Committee activities and progress are reviewed by the Chairman at these sessions. He also outlines future activities of the Committee.

One topic of special concern to all the vendors is their continued difficulty in obtaining high performance components from U.S. chip vendors. A large fraction of the IC's in a Cray computer are of Japanese manufacture. Fujitsu is the dominant supplier to Cray and, at the same time, Fujitsu is likely to become Cray's primary foreign competitor in the supercomputer marketplace. All three vendors claimed some evidence of Japanese companies not shipping their latest high performance components to outside markets until their needs for their own products are satisfied. This represents a serious problem to dependent U.S. companies who may be forced to use components that are a factor of two or more lower in performance than their Japanese competitors.

According to the U.S. supercomputer vendors, their component problem is caused by the fact that the supercomputer market offers a small number of chips compared to other markets such as microcomputers. The combination of a small market, together with high development costs and high capital equipment costs make the high performance component business unattractive to U.S. chip makers. The large, vertically integrated Japanese companies such as Fujitsu can justify development of these components based on their own product needs. Several Committee meetings with U.S. semiconductor manufacturers have confirmed this view of the problem.

Because of the seriousness of the component problem to the U.S. supercomputer industry, the Committee will continue to examine this problem with the intent of recommending solutions.

A similar problem exists with high performance peripherals. In particular, there continues to be concern over development of high performance rotating storage to meet the requirement of new supercomputer systems. One of the most encouraging developments is a Government (NASA and DOD) project on magneto-optical storage. Although not originally intended for supercomputer use, the storage density, data transfer rates, and price appear promising. The Committee has helped to make the supercomputer vendors aware of this development and one company, ETA Systems, is considering investing in the development.

## Interagency Sharing of Government-Owned Supercomputing Resources

The Committee is developing an arrangement among agencies with supercomputer centers to exchange services. The primary benefit will be to research groups that receive support and supercomputer time from more than one agency. The idea is to allow such a research group to obtain its supercomputing from one center rather than to be forced to use two or more centers with different operating systems or even different computers which would be inefficient. In addition, trading supercomputer time will allow contractors of the various agencies access to varied computational resources when appropriate. In order to provide a practical arrangement for the sharing of supercomputer facilities, an agreement that conforms to existing laws and directives is in preparation. Rules for the use of communications network will be established among users will be provided.

## Training for New University Supercomputer Centers And Users

The Committee has been concerned with training of both operators of new centers and new users. Most of the Nation's expertise in the operation and use of supercomputers resides in U.S. Government laboratories funded by DOE, NASA, and NSF (NCAR) and thus these laboratories represent the largest supercomputer training resource in the Nation. The Committee has taken steps to obtain the help of these laboratories to train operators of new supercomputer centers and several new centers are taking advantage of this opportunity. In addition, the Committee has been concerned with the need for training new users in the scientific community. As a result of Committee discussions, the National Science Foundation organized several workshops this past summer.

Network Subcommittee

With the significant increase in the number of Government supported supercomputer centers and the need to provide communications networks to serve the rapidly growing user community, the Committee recognized the need to address the communications problem across the Government and established a network subcommittee. This subcommittee will act to coordinate agency networking activities in order to avoid duplication of effort and to maximize the productivity of the Nation's research community. The subcommittee is currently adressing three topics:

1. Individual agency status, capabilities and plans are being documented and disseminated to other agencies.

2. A move toward a scientific network infrastructure is being examined which would tie together individual agency backbone networks via standard network functionality and gateways to link the scientific communities of individual agencies.

3. A plan for migration to the ISO standards is being addressed.

Facilitating Industry Access To Government Sponsored Research Programs

Certain Government sponsored R&D programs are developing technology that can contribute to advancing U.S. computer technology . These government programs generally represent an investment in high risk, long-term R&D that has a direct national defense purpose. Some of the results from such R&D can be used by the U.S. supercomputer industry in the relatively near-term. Two such programs are the USODRE funded Very High Speed Integrated Circuit (VHSIC) and DARPA funded gallium arsenide research. The Committee arranged to obtain access to these programs for the U.S. supercomputer vendors by having them invited to the annual program reviews. The vendors indicated that they made some very useful contacts in the semiconductor industry through these program reviews.

As mentioned above, the Committee hosted an interagency presentation of proprietary RCA and 3M magneto-optical technology with promise for meeting data transfer rate requirements and large peripheral storage needs of future generations of supercomputers. (The prospects for improving conventional disk drives to meet forecast requirements are not promising.) Subsequently, the developers held joint meetings to brief supercomputer manufacturers. Committee policy is to encourage commercial development of new technology by the private sector. Should situations arise where commercial exploitation of crucial supercomputer technology is not commercially feasible because of market limitations, the Committee will so report with recommendations for appropriate action.

## Interagency Information Exchange

The Committee regularly receives briefings from Government agencies. For example, NSA reported on the structure and mission of the newly organized Supercomputer Research Center in Maryland.

# FCCSET COMMITTEE MEMBERSHIP LIST

Committee Chairman:

Dr. James F. Decker
Deputy Director, Office of Energy Research
US Department of Energy
1000 Independence Avenue SW
Washington, DC 20585
(202/252-5434)

Members:

Dr. Paul B. Schneck
Director of the Supercomputing Research Center
4380 Forbes Boulevard
Lanham, MD 20706
(301/731-8065)

Dr. John Riganati
Supercomputing Research Center
4380 Forbes Boulevard
Lanham, MD 20706
(301/731-8065)

Dr. John Connolly
National Science Foundation
1800 G Street NW
Washington, DC 20550
(202/357-7558)

Dr. John McTague
Office of Science and Technology Policy
New Executive Office Building
Washington, DC 20506
(202/395-3961)

Mr. Kermith Speierman
National Security Agency
Ft. Meade, MD 20755
(301/688-6434)

Mr. James Burrows
National Bureau of Standards
The Technology Building
Gaithersburg, MD 20899
(301/921-3151)

Members (con't):

Dr. Craig Fields
DARPA/IPTO
1400 Wilson Boulevard
Arlington, VA 22209
(703/694-5922)

Dr. Randolph Graves
National Aeronautics and Space Administration
600 Maryland Avenue NW
Washington, DC 20546
(202/453-2828)

Dr. Allan T. Mense
Special Assistant to the Chief Scientist
Office of the Secretary of Defense
SDIO, The Pentagon
Washington, DC 20301-7100
(653-0131)

Executive Secretary:

Mr. Norman Kreisman
US Department of Energy
Office of Energy Research
Field Operations Management
1000 Independence Avenue SW
Washington, DC 20585
(202/252-9746)

File

# REPORT OF THE FEDERAL COORDINATING COUNCIL ON SCIENCE, ENGINEERING, AND TECHNOLOGY PANEL ON ADVANCED COMPUTER RESEARCH IN THE FEDERAL GOVERNMENT

JUNE 1985

# REPORT OF THE FEDERAL COORDINATING COUNCIL ON SCIENCE, ENGINEERING, AND TECHNOLOGY PANEL ON ADVANCED COMPUTER RESEARCH IN THE FEDERAL GOVERNMENT

JUNE 1985

## TABLE OF CONTENTS

## LIST OF FIGURES

## LIST OF TABLES

In the course of its initial discussions, the panel focused on the Federal role in advancing the state-of-the-art in computer science research. However, the sheer size and diversity of the field made categorization, much less coordination, a challenge. It also became apparent that coordination was most appropriate and desirable in the subset of research which the panel identified as Very High-Performance Computing; each of the participating organizations has considerable interest in that area, and the potential for joint funding and/or sharing of results is greatest. The Panel considered the mechanisms and strategies available to the Government to maximize its leverage and accomplish its stated goals and objectives. The panel also considered the appropriate Government responses to the "threat" that foreign programs may pose to the competitive and economic position of the United States. The FCCSET panel on research coordination plans to report periodically on Very High-Performance Computing throughout the Government and continue to meet regularly to maintain high-level coordination among Federal agencies and departments. As part of its ongoing activities report, the panel plans to assess the status and effectiveness of collaborative activities in this field.

The recommendations cited herein address key issues related to ensuring that the United States not only retains but advances its position of leadership in the Very High-Performance Computing field. Since the other FCCSET supercomputing panels have studied procurement strategies relative to the commercial market and provisions for greater access to supercomputers, these issues are not covered here.

In making its recommendations, the panel recognized that sustained and growing support for basic research is necessary for advancement in the Very High-Performance Computing field. Other related areas also need to be addressed in order to successfully move existing technology out of the laboratory and into applications. Overall, these recommendations are intended to ensure that the United States continues to be the prime source of Very High-Performance Computing technology in the decades ahead.

# EXECUTIVE SUMMARY

The United States has long been recognized as the leader in the computer and information processing field and despite the challenges presented by foreign initiatives, has retained its leadership due to the strong technical, industrial, and academic base and research support that have evolved over the past decade. Continuing efforts within the research community, however, have not prevented this position from eroding rapidly. Although the United States remains dominant in information processing, the Japanese and others have targeted both high-end numerical computing, or supercomputers, and symbolic processing for technological exploitation on their quest to capture leadership in the information processing field.

Since the early 1960's, the key technical and programmatic issues in high-performance computing have been investigated by a series of studies, reports, and workshops sponsored primarily by the Federal Government. Interest in these issues declined in the mid-1970's as Government funding levels fell off, but has been revived by the perceived threat of foreign competition, the emergence of innovative architectural concepts in parallel and multiprocessor machines for scientific and symbolic computation, and the identification of a broad spectrum of mission-specific applications that will require radical improvements in the speed and performance of computing systems. Renewed Government interest and funding have spawned such activities as DARPA's Strategic Computing Program, NSF's Advanced Scientific Computing program, DOE's Energy Science Supercomputing Program, and DoD's Supercomputer Research Center. In FY85, the total Federal investment in Very High-Performance Computing research will approach $101 million.

To address the need for the next generation of both numerical and symbolic computing capabilities, the Federal program in Very High-Performance Computing encompasses a wide spectrum of activities in Advanced Computer research and development. Significant high-performance parallel and multiprocessor computing advances will come from research efforts in machine architecture, advanced programming languages, and systems and software methodologies for concurrent multiprocessor operations. Areas such as system software for parallelism--including artificial intelligence techniques applied to debugging, testing, verification, and performance measurement--will require significant investment over the next decade.

Researchers in high-performance numerical computing and symbolic processing expect significant gains in performance from advanced multiprocessor architectures. Although the design and development of multiprocessors for symbolic and numerical problems are now being explored separately for technical reasons, significant similarities in the underlying architecture concepts may emerge as basic research efforts continue. Certain numerical machines are capable of performing symbolic computations quite rapidly, but they have not been optimized for this class of applications and generally are regarded as not cost-effective. Existing high-performance computers for numerical and symbolic domains are based widely on varying architectural concepts, but long-term trends indicate that important common, underlying aspects, with complementary domain-specific specializations, will result. It should be possible to develop supporting hardware and systems software to allow multiprocessors to operate in both symbolic and numeric environments concurrently. Further, for these systems, there is a critical need for characterizing and measuring computer performance to permit designers and users to discriminate among alternative architectures.

Federal high-performance computing research consists of a number of programs and activities carried out by individual Government agencies and by combinations of Government agencies working together. This overall effort supports the Government's mission to:

(1) Conduct basic and applied research in computational sciences and engineering, specifically in high-performance computing;

(2) Develop promising concepts into prototype systems where appropriate;

(3) Evaluate the performance of existing and planned high-performance computing technology;

(4) Ensure the necessary infrastructure for the conduct of the program; and

(5) Apply high-performance computing in meeting mission-specific goals.

Coordination among the Government agencies and departments responsible for funding advanced computer science research is carried out at the program manager level. Program managers are responsible for sharing research plans and results; reaching agreement on generic and mission-specific goals and on program interactions and modifications; discovering relationships, overlaps, gaps, and opportunities in various research activities; and creating a coherent total-program approach. Program managers throughout the Government work to get the most from their funds by sponsoring programs in a joint or complementary manner where appropriate. Although it is impossible to quantify the amount of formal or informal coordination taking place, it is substantial. Program managers routinely refer research proposals to one another if certain tasks, or the entire proposal, seem more appropriate to the other agency's missions and programs.

By providing a forum for the exchange of information about individual Government agencies' research and development programs in very high-performance numerical and symbolic computing, this panel has developed a global perspective on interagency coordination. The existing coordination process could be more visible and the panel will continue to emphasize the visibility issue as the support for very high-performance computing research and development in the Government increases.

In the course of its deliberations, the panel developed a set of recommendations to enhance the United States position as the leader in very high-performance computing. The recommendations are summarized as follows:

- Maintain a vigorous, coordinated research program.
- Increase emphasis on understanding fundamental issues in parallel processing.
- Promote research activities that apply to a broad class of problems.
- Do not over-coordinate basic research.
- Explore a diverse set of architectures.
- Coordinate exploratory machine architecture development efforts.
- Improve technology transfer mechanisms from Federally sponsored research to the commercial sector.
- Develop programs designed to augment the number of trained researchers.
- Take steps to ensure that compensation is adequate to retain qualified researchers in the public and academic sectors.
- Develop effective performance measurement and modeling techniques.
- Investigate the infrastructure requirements to support the research community.
- Maintain a visible interagency coordination effort.

The FCCSET Panel on Advanced Computer Research considered the impact of the Federal research program in advancing the United States capabilities in very high-performance computing, and the key support issues such as facilities, personnel, and Government investment in the requisite technologies. The recommendations that have been developed provide a basis for the continued development of the technical expertise the United States requires to maintain its leadership position in information processing and advanced, very high-performance computer technology.

## MEMBERS OF THE FCCSET PANEL ON ADVANCED COMPUTER RESEARCH IN THE FEDERAL GOVERNMENT

Dr. Robert E. Kahn, Defense Advanced Research Projects Agency, Chairman

Dr. Donald Austin, Department of Energy

Mr. Ross Bainbridge, Department of Commerce, National Bureau of Standards

Dr. Bernard Chern, National Science Foundation

Dr. Ronald L. Larsen, National Aeronautics and Space Administration

Dr. Joseph Markowitz, Central Intelligence Agency

Dr. K. M. Speierman, National Security Agency

Dr. Andrew Pettifor, Office of Science and Technology Policy

Additional contributors to the report are listed in Appendix E.

## PREFACE

In the spring of 1983, the Office of Science and Technology Policy (OSTP) formed three interagency panels under the Federal Coordinating Council on Science, Engineering, and Technology (FCCSET) to examine critical emerging issues in the computer field. Two of the FCCSET panels were chaired by the Department of Energy (DOE); these panels focused on the issues of procurement and access to high-speed numerical machines known as supercomputers. Reports covering these two areas were prepared for OSTP by the end of 1983 and distributed widely throughout the government. The two panels were later merged into one to consider the follow-on issues in supercomputer procurement and access. The third FCCSET panel is chaired by the Department of Defense (DOD) and its charter is to stimulate the exchange of information within the government on high-performance symbolic computing and artificial intelligence (AI).

During the course of its initial deliberations, it became clear that interest in information exchange extended beyond symbolic processing and AI. Questions arose concerning the architectural and computational differences between symbolic and numerical computing and the requirements of very high-performance machines that might be able to handle both types of computations effectively. As the individual agencies presented their current and planned research programs, the relationship between these two areas of research became a topic of discussion in its own right. Consequently, it was decided to broaden the charter of the panel beyond symbolic processing and artificial intelligence and to include federal research efforts in very high-performance scientific and numerical computing and Advanced Computer Research in general.

This document constitutes the first report of the FCCSET Computer Research Coordination Panel. It provides an overview of current and planned federally sponsored research activities in Advanced Computer Research and Very High-Performance Computing, in particular. It summarizes existing activities funded by individual agencies and by more than one agency, identifies the FY 1983 to FY 1985 funding, and presents a set of findings and recommendations for further consideration.

Within the Government, the primary supporters of basic computer science research are the National Science Foundation (NSF), Department of Energy (DOE), and the Defense Advanced Research Projects Agency (DARPA). The National Aeronautics and Space Administration (NASA), and the military services, while conducting limited basic research programs, complement these efforts with mission-specific investigations. In addition, the Department of Commerce, through the National Bureau of Standards (DOC/NBS), is concentrating on standards, metrics, and benchmarks, while the National Security Agency (NSA) and the Central Intelligence Agency (CIA) are studying the intelligence aspects of the technology.

# REPORT OF THE FEDERAL COORDINATING COUNCIL ON SCIENCE, ENGINEERING AND TECHNOLOGY PANEL ON ADVANCED COMPUTER RESEARCH IN THE FEDERAL GOVERNMENT

## A. INTRODUCTION

The United States has long been recognized as the leader in the information processing field and, despite the challenges presented by foreign initiatives, it remains in that position as a result of its strong technical industrial base and continuing research support. In spite of vigorous continuing efforts within the research community, however, this lead is being threatened. Foreign programs aimed at Very High-Performance Computing signal the need for a vigorous national response by the United States to maintain its leadership.

Since the early 1960's, the key technical and programmatic issues in Very High-Performance Computing have been documented in a series of studies, reports, and workshops sponsored primarily by the Federal Government (Lax, et. al., 1982; Feigenbaum, et al., 1983; Schwartz, et al., 1984; and Decker, et al., 1984). Renewed Government interest and funding have spawned such activities as DARPA's Strategic Computing Program, NSF's Advanced Scientific Computing Program, DOE's Energy Science Supercomputing Program, and DoD's recently announced Supercomputer Research Center. Research in Very High-Performance Computing is not yet a major segment of the overall Federal effort in Advanced Computer Science research but it is clearly growing.

The Federal Coordinating Council on Science, Engineering, and Technology Panel on Advanced Computer Research in the Federal Government was originally tasked by the Office of Science and Technology Policy to stimulate the exchange of information within the Government on symbolic computing and artificial intelligence (AI). Subsequently, the panel broadened its charter beyond symbolic processing and artificial intelligence to include research in very high-performance scientific and numerical computing, with the panel focusing on the Federal role in advancing the state of the art in this field.

The purpose of this report is to provide a concise summary of ongoing Federal research efforts in Very High-Performance Computing, and to identify a set of findings and recommendations which we believe will strengthen our nation's overall capability in this extremely important area. This panel compiled a series of interrelated findings and recommendations on the issues that Government policymakers should consider in order to ensure that the United States retains its strong national capability in advanced computing technologies. Generally, the size and diversity of the many efforts within the Government, industry, and academic community result in special problems in technology transfer, availability of qualified personnel and necessary research facilities, and information flow. The recommendations take into account the importance of basic research as well as a number of key engineering issues that must be considered before technologies move from the research laboratory environment into mainstream applications. The recommendations, taken as whole, will assure that the United States continues to be the prime source of Very High-Performance Computing technologies in the decades ahead.

The Panel defines Very High-Performance Computing to mean efforts specifically concerned with the exploitation of concurrency and parallel processing to achieve dramatic increases in speed of computation. The main focus is on development of multiprocessor systems with emphasis on scalable architectures whose performance increases are near linear as more processors are added. Specifically included are computational models and architectures in parallel processes; problem decomposition techniques and languages for expressing parallelism; compilers and operating systems specifically devised for multiprocessor systems and software which manages and controls them; algorithms and heuristics designed for parallel processing; applications which will lead to an increased understanding of the underlying principles of these technologies; and those methods for quantitatively characterizing and measuring advances in each of these areas.

The Government also supports a significant amount of Advanced Computer Research which is targeted at improved functionality, reliability, etc. and which the Panel has excluded from its interpretation of Very High-Performance Computing. Most computer research that leads to incremental improvement in computation speed has been excluded from this category. Most artificial intelligence research is expected to contribute major functional advances in computation, but little of it is yet addressing performance speed-up and is not included. Further, most ongoing research in basic computer science and computational mathematics, such as efforts in software technology, distributed systems, numerical analysis, and computer networking which are aimed at substantial advances in functionality, are excluded. In addition, related information processing research efforts such as software life cycle maintenance, system reliability and microelectronics technology, although vitally important in development and applications, have not been included in either category. Approximately $298 million is being spent in FY85 on Advanced Computer Research, of which approximately 34 percent or $101 million is being spent on Very High-Performance Computing research.

The Panel observed that the approaches being taken for high-performance numerical computing and for symbolic processing are noticeably different at this time, yet researchers in both areas expect significant gains in performance from multiprocessing. Existing high-performance computers for the numerical and symbolic domains vary widely, but long-term trends indicate important common architectural aspects with complementary domain-specific specializations. Certain numerical machines are capable of performing symbolic computations quite rapidly, but they have not been optimized for this class of applications and generally are not regarded as cost-effective. Although the conceptual design and machine development of multiprocessors for symbolic and numerical problems are now being explored separately for technical reasons, significant similarities in the underlying architectures are expected to emerge later. The trend in the use of highly parallel, multiprocessor computer architectures toward combined numeric and symbolic applications may lead to architectures capable of executing both forms of computation in an efficient and cost-effective manner.

Advanced numerical computation addresses key generic areas such as:

1. Numerical analysis and simulation applied to computational models for which closed-form solutions do not exist;

2. Manipulation of very large volumes of data, perhaps generated from the numeric solution to the model analysis as in item 1; and

3. Graphic presentation of complex, multi-dimensional data.

Symbolic processing manipulates non-numeric objects and is concerned with issues such as:

1. Knowledge representation and semantic retrieval techniques;

2. Feature or symbol extraction, as in the translation of a signal into a symbolic representation;

3. Expert systems which combine knowledge with inferencing mechanisms; and

4. Search techniques based on non-exhaustive heuristic methods.

In certain cases, such as image understanding and speech recognition, numerical techniques are first used to prepare data for input to a symbolic processor, such as in generating a primal sketch, which consists of significant lines and edges in an image.

The Federal research effort in Very High-Performance Computing consists of a number of programs carried out by individual Government agencies or by combinations of Government agencies working together to meet specific national mission and application needs. Requirements for defense, space technology, energy technology, excellence in scientific research, as well as others, are reflected in the overall programmatic goals.

In the course of its deliberations, the panel focused on four key issues concerning the ability of the Government to conduct a successful basic research and exploratory development program in Very High-Performance Computing. These issues are:

1. **TECHNOLOGY** - What research areas must be exploited to advance the state-of-the-art;

2. **PERSONNEL** - How to ensure a sufficient number of researchers by the end of the decade, how will the projected shortfall be met in high-performance computing;

3. **FUNDING** - What is the current Government investment and how is it apportioned; and

4. **RELATED ISSUES** - Coordination of Federal activities, socio-economic impacts, and facilities to directly support research activities in high-performance computing.

The socio-economic environment for computer research and development in the United States is undergoing significant changes in order to meet the increasing demands of industry, Government, and the scientific community for numerical and symbolic computational capabilities. The recent formation of such entities as the Semiconductor Research Corporation (SRC), the Microelectronics and Computer Technology Corporation (MCC), the Parallel Processing Research Council, the recently announced Software Productivity Consortium (SPC), as well as the growing number of direct industry-academic partnerships in computer research and development, is in itself a benchmark of the concern being expressed by all sectors of the research community as to the strength of the United States' technical leadership position. The private investment in all categories of university research increased by 13 percent from 1982 to 1983 and exceeded $390 million. This private investment--in the form of fellowships, grants, or direct contracts--can be expected to continue to increase, especially as the programs within the joint ventures reach projected funding levels. The recently signed National Cooperative Research Act of 1984 facilitating the formation of joint research and development ventures will further promote private investment into the academic and industrial research community. On the other hand, given the charter of these joint ventures, it is not clear how the transfer of technology from the sponsored research to the commercial marketplace beyond a consortium's individual members will take place. In addition, if researchers are drawn from the academic community into projects conducted within the consortium's own staff, a decline in the ability to educate new graduate students may result as the number of qualified educators decreases.

The overall Federal program on Very High-Performance Computing research encompasses a wide spectrum of activities. Basic research and exploratory development programs in this area will lead to a new generation of very high-speed computers and provide the underlying theory and knowledge that will create an environment in which innovative ideas in information processing technology will flourish. These efforts will also be a key source of scientists and engineers required for the future growth in this rapidly expanding and critical field.

## B. THE FEDERAL ROLE

The various Government agencies and organizations involved in Advanced Computer Research support both generic and mission-specific programs. This report focuses on the Federal research activities in Very High-Performance Computing **with emphasis on those that are generic in nature** and which yield fundamental concepts, technology, knowledge, people and ideas. Activities directed toward an individual organization's mission-specific goals have not been included in determining the amount of funding for Very High-Performance Computing research.

Research and development can be viewed as a four-part process. The first step in this process, basic research, consists of fundamental undirected exploration at the conceptual level. Research into these concepts is often initiated by the investigator doing the research utilizing existing facilities, computer resources, and personnel, funded at low cost, and accomplished with little or no special equipment; coordination on a project-by-project basis is appropriate and effective. The key to the success of basic research programs is highly dependent upon the contributions of individual researchers. Federal coordination typically takes the form of peer review within the relevant Federal agencies. A researcher who has proposed a program may even be presented with a "coordinated" Government response to his or her project, which in turn can lead to funding of selected tasks of a total proposed program by different agencies. The specific areas selected for funding reflect the overall mission specific interests of each agency. However, this form of coordination encourages the entrepreneurial aspects of the research community to propose new areas for exploration, even if they go beyond the individual interests of the single agency or department of the Government.

The second part of the process, exploratory development leading to experimental capabilities, is concerned with the feasibility of applying promising basic research results, often to generic classes of problems. In the hardware domain, a prototype or "breadboard" might be constructed along with elementary system software. At this point research begins to require increasing capital resources and facilities; the need for more extensive coordination begins to emerge. Multiple parallel efforts, although often desirable at the basic research level, may be unaffordable for exploratory development, so refinement and selectivity may be necessary here. On occasion, several sponsors may pool their resources toward a common goal, or each may support complementary aspects of a program. Examples include the Cosmic Cube (architecture funded by DARPA, applications by DOE), the Wave-front Array Processor (basic research funded by NSF, signal processing application by ONR), and Systolic Array Processors (computational model and algorithm research funded by NSF and ONR, architecture and breadboard by DARPA).

The third and fourth phases of the process involve advanced development and production engineering. At these levels, program costs are far greater and only a few efforts can typically be pursued. The selection of candidates for advanced development is done individually by each organization in response to mission-specific needs. The need for coordination across organizational lines rapidly decreases as the technologies move into these latter phases.

The Federal Program structure in Advanced Computer Research is depicted in Figure 1 as a series of building blocks or supporting layers of which the research program in Very High-Performance Computing is a part. As indicated above, **emphasis was placed on programs that are generic in nature. Mission-specific programs were not included in determining the overall Federal funding profile.** Also, programs related directly to purchasing computers or facilities for other scientific and engineering disciplines were not included in any of the research figures, an example being the recently announced NSF program on advanced scientific computing.

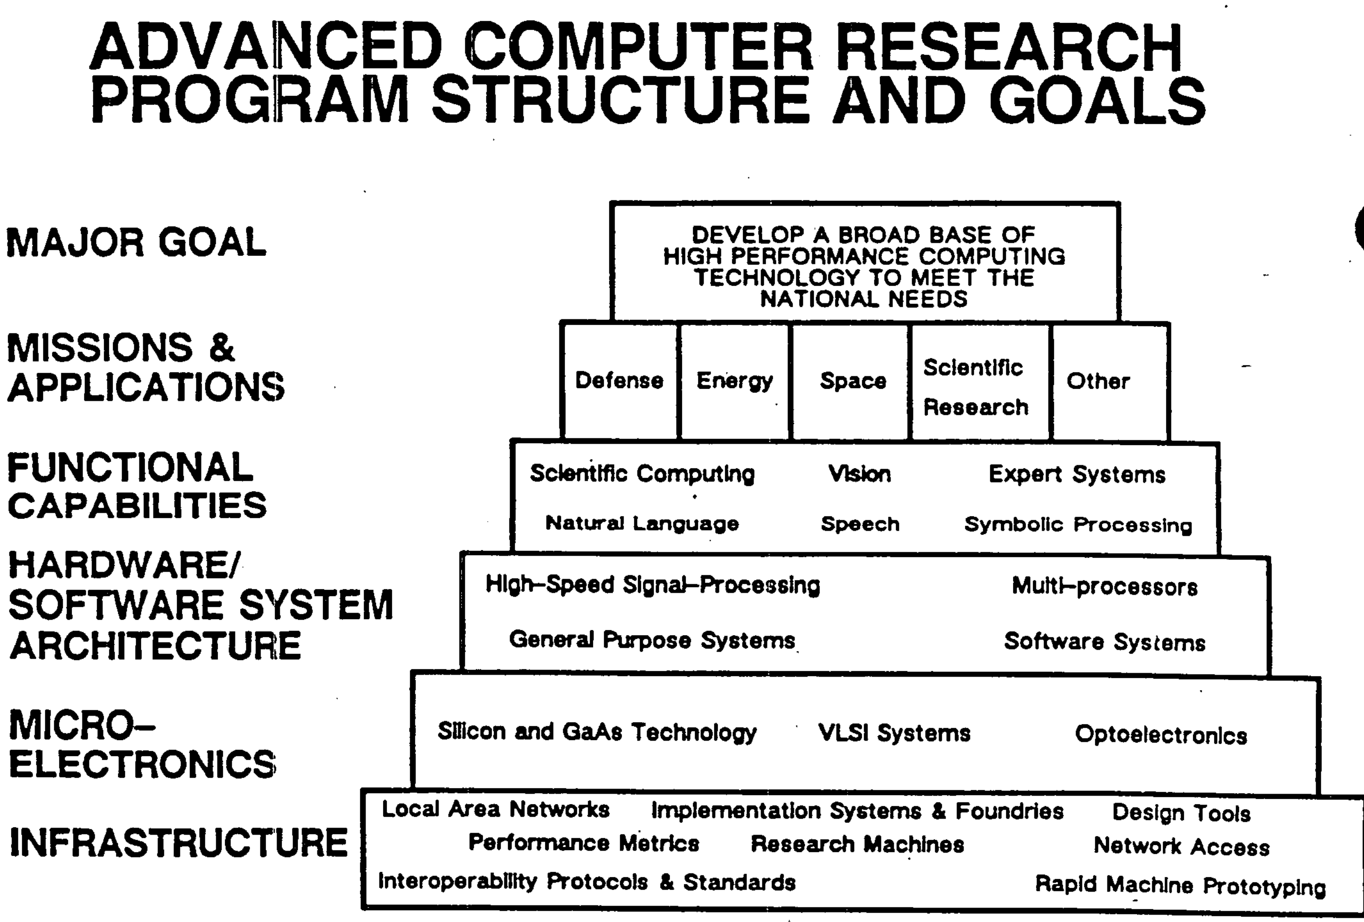


Figure 1

Federal agencies interact in a variety of ways to coordinate their efforts, including:

1. Seeking agreement on generic and mission-specific goals;

2. Sharing research plans and results;

3. Discovering relationships, potentially overlapping interests, and opportunities in various research projects;

4. Conducting joint research programs; and

5. Creating a coherent, combined research program encompassing Government, industry, and academic sectors while maintaining mission-specific goals of individual organizations.

For many years, coordination among the agencies and departments responsible for funding computer science research has been effectively carried out at the program manager level. More recently, the FCCSET panels have provided a higher level of coordination as well. Although it is impossible to quantify the existing formal, semiformal, and informal interagency coordination, and although it is not completely visible outside the participating organizations, the amount of coordination is substantial. The authors of this report recognize the importance of providing increased visibility to this process as the support for Very High-Performance Computing grows. The report of this panel is expected to assist the coordination process by summarizing the ongoing and planned high-performance computing research and development activities in the Government. The panel itself is the most visible example of inter-agency coordination and visibility into the coordination process will be enhanced by the regular reviews of the FCCSET panel.

C. DEFINITION OF CATEGORIES

Advanced computer research is comprised of a set of multidisciplinary basic research and exploratory development activities, for which the panel developed the following set of nine categories to characterize them:

1. **Computational Mathematics:** The design, analysis, and implementation of algorithms for solving basic numerical problems by computer.

2. **Computer Architecture (Hardware and Software):** The design, simulation, and development of new computer architectures including both hardware and system software.

3. **Machine Intelligence and Robotics:** The development of software and conceptual designs which allow computers to carry out tasks which would be considered "intelligent" if performed by a human.

4. **Distributed Computing and Software Systems:** Techniques and procedures for building systems consisting of multiple computers connected by communication networks; the technology for designing and building software systems, including rapid prototyping and reliable operation.

5. **VLSI Design and Special Purpose Computing:** Tools and techniques for designing state-of-the-art VLSI, and the development of innovative circuits and computing systems using those tools and techniques.

6. **Data Management:** Design and development of advanced data base management concepts and systems.

7. **Theoretical Computer Science:** The analytic study of fundamental problems in Computer Science.

8. **Network and Research Facilities:** The provision of facilities for communication and computation specifically for the purpose of furthering advanced computer science research.

9. **Performance Evaluation and Modeling:** The analysis and study of performance evaluation, metrics and models, operational standards and benchmarks, user interface technologies, and human factors.

These research areas are not entirely separate and distinct but interdependent. For example, studies in theoretical computer science may lead to the understanding of techniques to be applied by computational mathematics, which in turn can lead to new approaches or tools for constructing highly complex VLSI designs for multiprocessor architectures. Network and research facilities generally support the work in the other categories.

Figure 2 shows Very High-Performance Computing depicted as a small part of Advanced Computer Research which, in turn, is a small part of Information Processing research and development. The definition of Advanced Computer Research is itself subjective. The shaded area in Figure 2 indicates the panel's qualitative assessment of the coverage of the research activities, and thus the funding profiles, in this report under Advanced Computer Research. Programs may span across the boundary between the regions. In identifying programs and categorizing the relevant research activities, the result may omit some efforts not totally within Advanced Computer Research or may include those elements of Advanced Computer Research programs that are outside of the Advanced Computer Research arena. For example, the DoD STARS and ADA programs were judged part of the larger information processing category, as was much of the computer network research. Program application or mission specific computer research activities were specifically not included in this report. In addition, none of the proposed SDI efforts were included since at the time this report was prepared, that program was still being defined. In developing the funding profiles for Advanced Computer Research and Very High-Performance Computing Research, we were unable to insure systematically that all such funds were properly categorized. We believe the funding data presented in this report represent quite close estimates, based upon the information available to the Panel. In most cases, research efforts were decomposed into several tasks in order to allocate the funding by the categories identified by the Panel. As such, the specific funding breakdown presented may not reflect actual categories of funding budgeted by the various organizations.

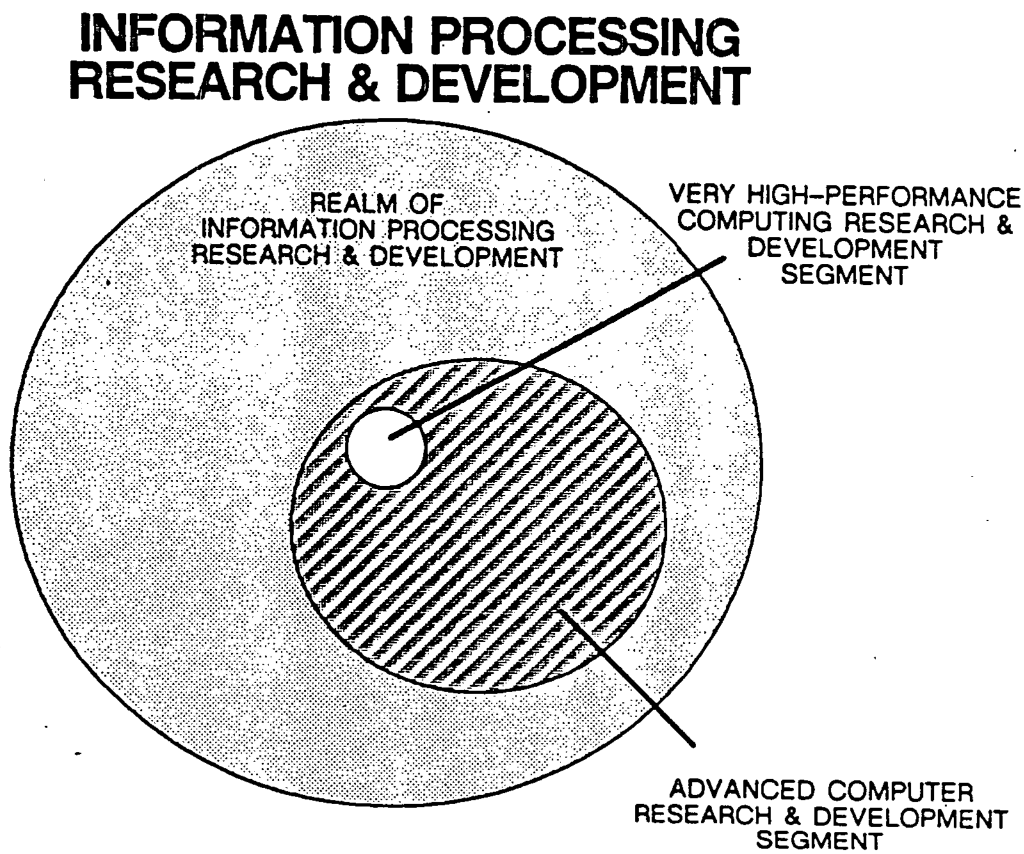


Figure 2

The panel compiled funding numbers for both Advanced Computer Research and for Very High-Performance Computing but focused its coordination efforts on the latter. Parallel computing architectures, many of which are still highly experimental, promise large increases in computational power without significant redesign of their components--that is, they are "scalable." Such architectures may make possible machines that are several orders of magnitude faster than the fastest existing machines. Scalable parallel architectures have major potential advantages in a VLSI environment because they can be built from large numbers of identical parts that can be mass-produced efficiently, because the design cost can be amortized over multiple configurations differing only in size, and because they are cost-effective.

The development of many small scale parallel experimental machines is seen as a necessity to explore all the promising ideas in multiprocessor architectures. Although device speeds continue to increase as the minimum feature size decreases, the physical limits of MOS technology will soon be reached. Multiprocessors offer the possibility of dramatic increases in speed through parallelism. There is a critical need for characterizing, measuring and modeling computer performance to permit designers and users to discriminate among alternative architectures. The speed of technology development in industry is increasing, and we are beginning to witness a new class of mini-supercomputers emerge in the marketplace. By the end of the decade, we expect to see the power of today's fastest supercomputers at the price of today's most powerful minicomputers.

## D. FEDERAL ACTIVITIES AND FUNDING

The Federal investment for FY 1983, 1984, and 1985 in Advanced Computer and Very High-Performance Computing Research is shown in Figures 3 and 4, respectively, by the major funding organizations. Individual organization funding is summarized in Appendix A, Tables A-1 and A-2. To further delineate the Federal investment in these two areas, the panel reviewed the Federal efforts in each of the nine categories in Section C. The expenditures in these research categories in FY 1983, FY 1984, and FY 1985 are shown for Advanced Computer and Very High-Performance Computing Research in Figures 5 and 6, respecitvely, with detailed funding by area given in Appendix A, Table A-3. The total funding by organization and research category was compared, and the results are summarized below.

It has been difficult to develop precise funding figures for both Advanced Computer Research and Very High-Performance Computing since most Federal programs are not fiscally structured along those lines. The figures listed herein reflect the best estimates of spending in each of the categories.

The Federal investment in Very High-Performance Computing Research was approximately $36.6 million in FY 1983, $57.8 million in FY 1984 and $100.9 million FY 1985 as shown in Table A-2. Figure 7 graphically shows the relative Federal investment in the Very High-Performance Computing component as a part of the overall Advanced Computing Research activity. Very High-Performance Computing Research comprised 21.1, 25.5, and 33.8 percent of the Advanced Computing Research funding in FY 1983, FY 1984, and FY 1985, respectively. The total growth rate for Very High-Performance Computing Research funding has exceeded by approximately a factor of 1.6 the growth rate in the overall Advanced Computer Research program from FY 1983 to FY 1985.

The university community forms a significant portion of the resources that are critical to the success of the Federal program in Very High-Performance Computing Research. Tables A-1 and A-2 present, for comparison, the Federal funding in Advanced and Very High-Performance

# FEDERAL INVESTMENT IN ADVANCED COMPUTER RESEARCH BY ORGANIZATION, FY 1983–1985

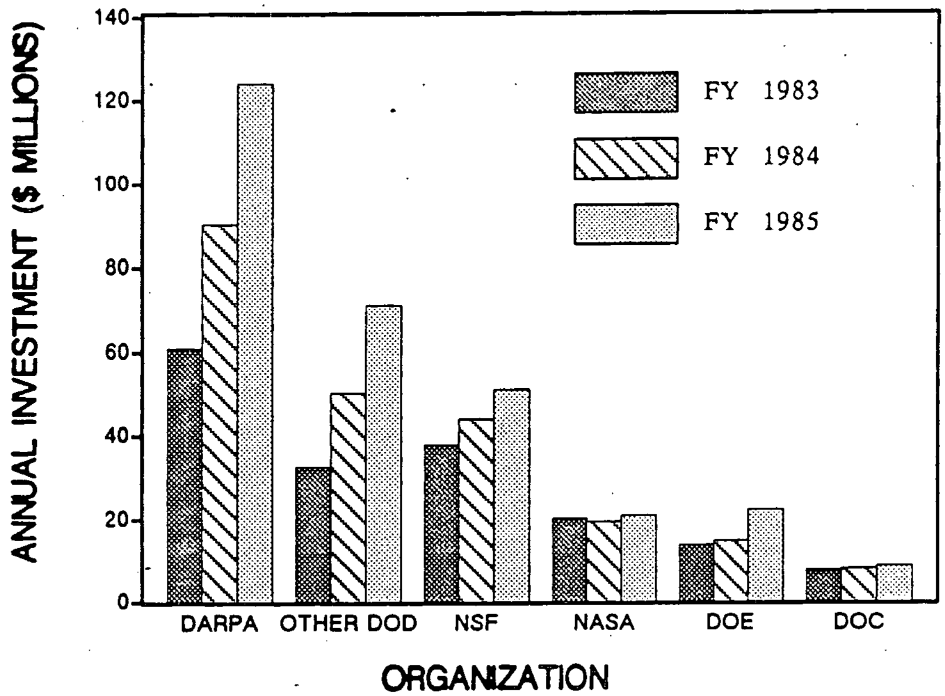


Figure 3

# FEDERAL INVESTMENT IN VERY HIGH–PERFORMANCE COMPUTER RESEARCH BY ORGANIZATION, FY 1983–1985

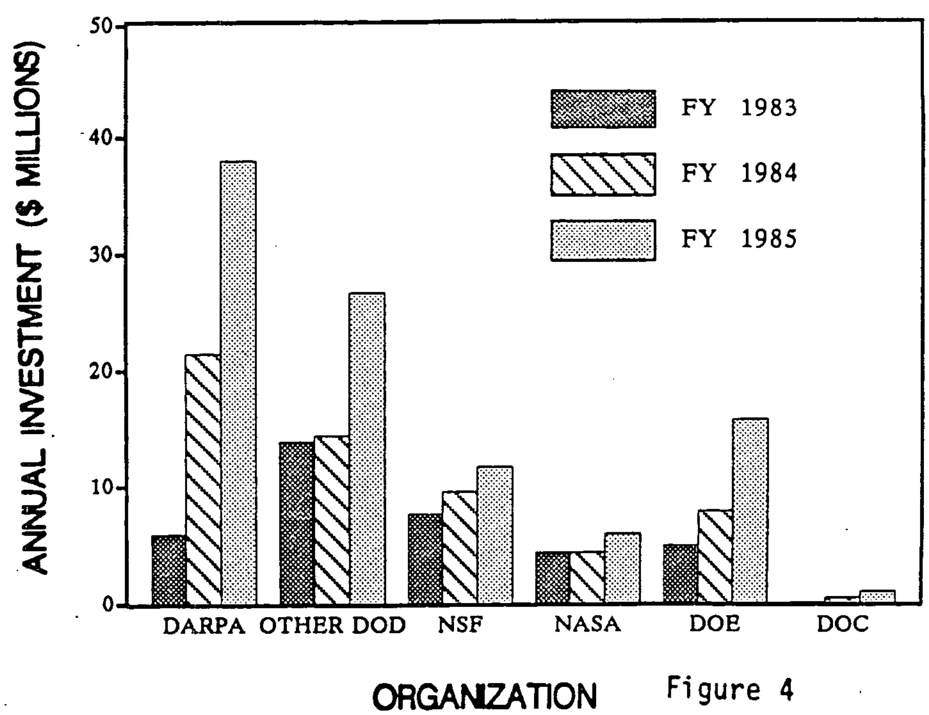


Figure 4

## FEDERAL INVESTMENT IN ADVANCED COMPUTER RESEARCH BY RESEARCH AREA, FY 1983–1985

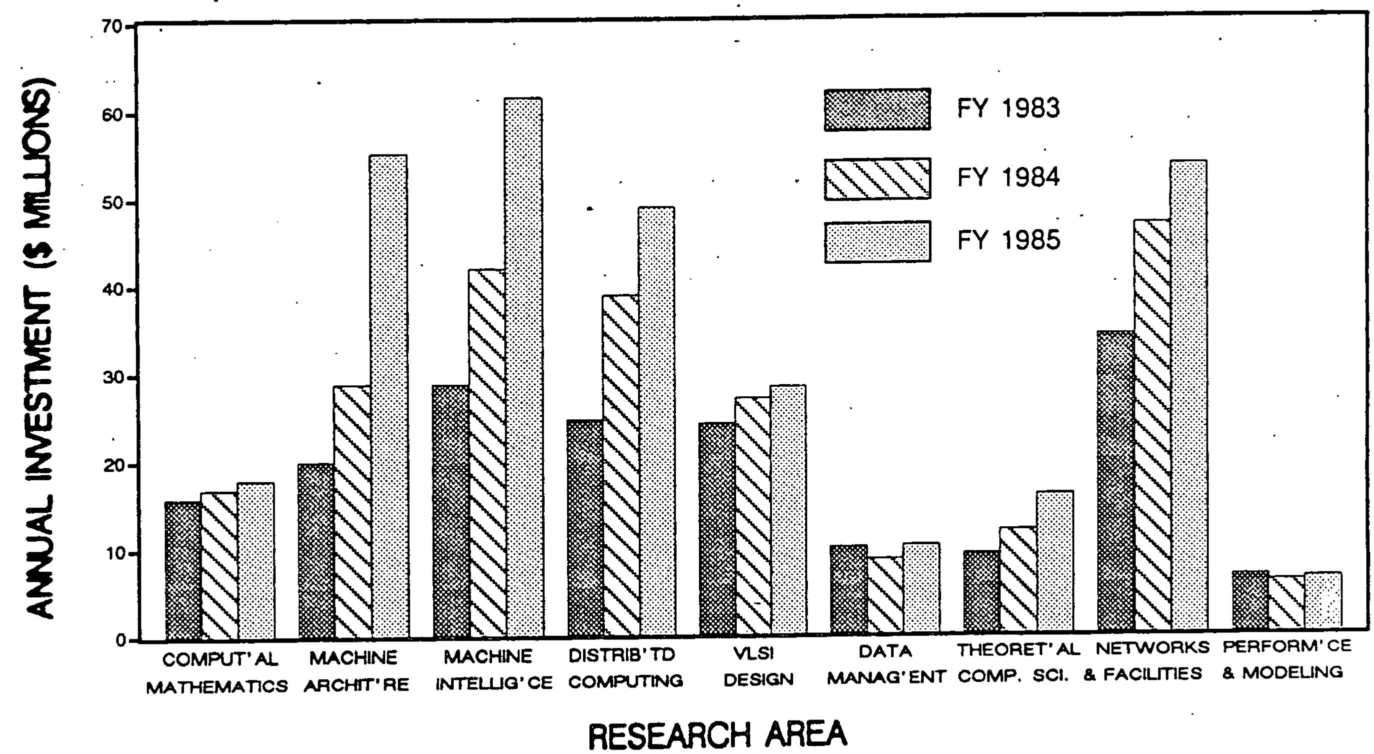


Figure 5

## FEDERAL INVESTMENT IN VERY HIGH–PERFORMANCE COMPUTING RESEARCH BY RESEARCH AREA, FY 1983–1985

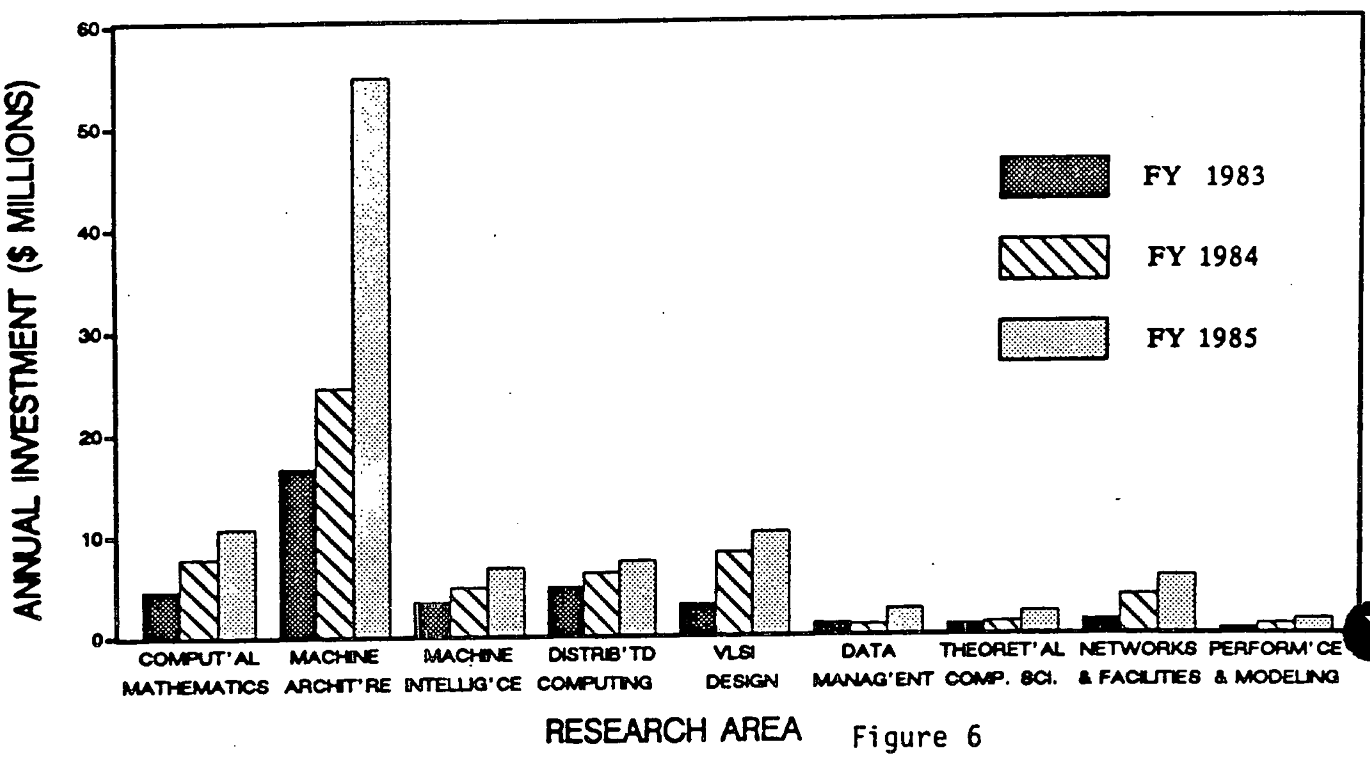


Figure 6

Computing Research, respectively, that is directed to the university sector. In FY 1985, it is projected that university funding will total $180.0 million, comprising approximately 60.3 percent of the budget for Advanced Computer Research, and approach $55.9 million, or 55.4 percent, of the Very High-Performance Computing Research budget. The university community has seen an increase from FY 1983 to FY 1985 of $67.5 million in its annual funding in the overall Advanced Computer Research program and an increase of $30.8 million in its annual budget in Very High-Performance Computing Research. Figure 7 also presents the relationship of the university segment to the overall funding profiles in this field.

## FEDERAL INVESTMENT IN ADVANCED AND VERY HIGH-PERFORMANCE COMPUTING RESEARCH: TOTAL PROGRAM AND UNIVERSITY COMPONENT

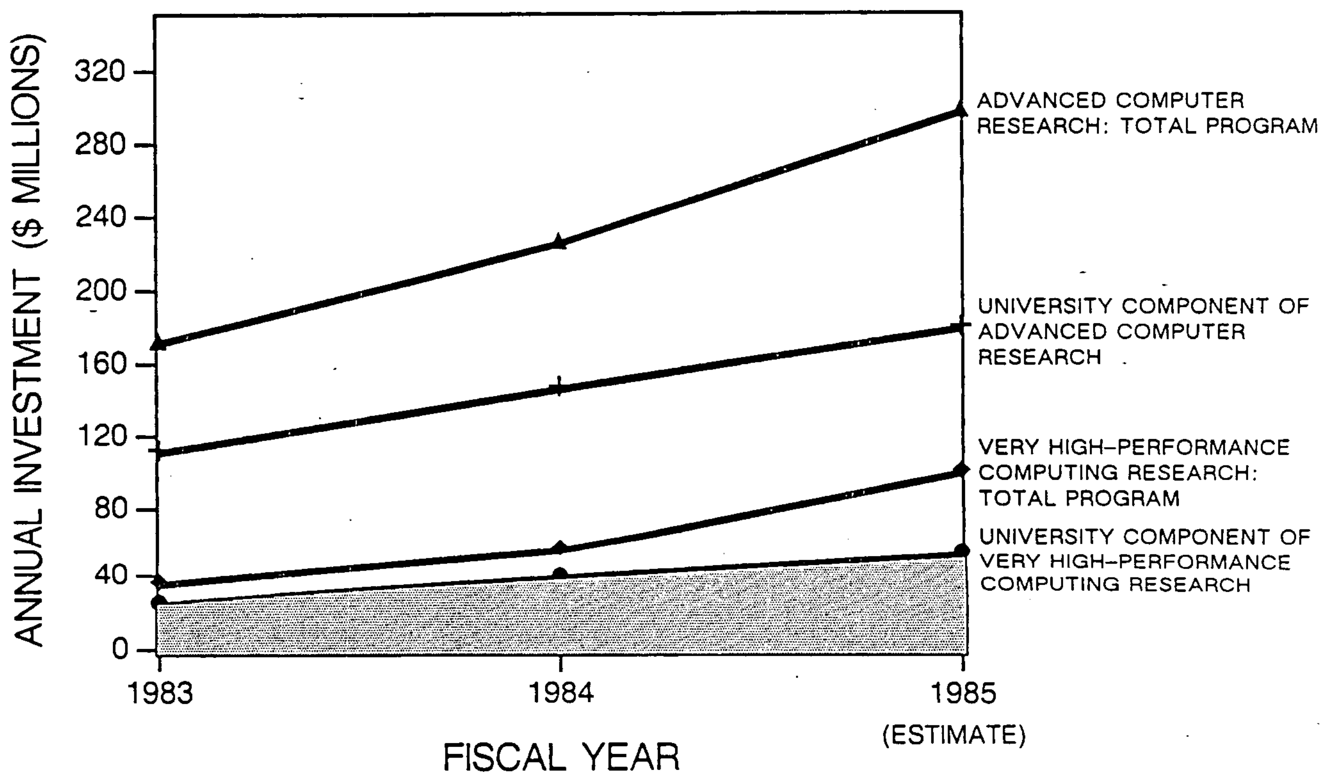


Figure 7

The funding of research in the area of machine architectures has exhibited the largest absolute growth of any of the nine categories. This growth is a reflection of the importance and the challenge that now confronts present day machines to meet the information processing requirements demanded by both civil and military systems. Specifically, in the multiprocessor area, several dozen studies and small scale explorations are underway. The Schwartz report (J. T. Schwartz, et al., "Report of the Research Briefing Panel on Computer Architecture," 1984) noted the tradeoff between the number of processors comprising a computer architecture and their size and complexity (or "granularity"). A coarse-grain design uses a relatively small number of high-speed processors, typically a microprocessor available as a commercial product. A fine-grain design strives for as much parallelism as possible with a large number of small, but carefully designed elementary processors each of which is fast but low in overall complexity. In general, these machines can execute multiple instructions simultaneously depending on their architecture. They may have a single memory address space shared by all processors or a separate memory address space per processor.

In FY 1984, Federal agencies were conducting research into thirty coarse-grain and eight fine-grain architectures. The performing organization and the Federal sponsoring agencies are listed in Appendix D, Tables D-1 and D-2 for coarse-grain and fine-grain architectures, respectively. A summary of the Federal investment in coarse- and fine-grain multiprocessor architecture research is shown in Table 1. The investment totals for fine- and coarse-grain architecture programs shown in Table 1 represent a significant portion of the total investment in machine architecture. In general, the fine- and coarse-grain totals reflect experimental prototype machines, simulation and exploratory development, whereas the remainder of the funding as shown in Table A-3 addresses conceptual architecture research.

TABLE 1

**FEDERAL INVESTMENT IN ADVANCED MULTIPROCESSOR ARCHITECTURE DESIGNS**

(FY 1984, $ MILLIONS)

| | COARSE-GRAIN DESIGNS | | FINE-GRAIN DESIGNS | |
|---|---|---|---|---|
| AGENCY | NUMBER OF PROGRAMS | FUNDING | NUMBER OF PROGRAMS | FUNDING |
| DARPA | 5 | $5.3 | 4 | $4.5 |
| AIR FORCE | 2 | 0.4 | | |
| NAVY | 4 | 3.5 | 1 | 0.4 |
| NSF | 17 | 4.4 | 3 | 0.4 |
| DOE | 9 | 2.6 | | |
| NASA | 7 | 0.9 | 1 | 1.0 |
| TOTAL | 30* | $17.1 | 8* | $6.3 |

* May not equal column total due to joint programs.

The Federal funding data for the fiscal years 1983, 1984, and 1985 for Advanced Computer Research by organization is detailed in Appendix B, while Appendix C presents similar data for Very High-Performance Computing Research, in accordance with the nine research categories described in Section C. From these tables, areas of significance to each agency can be seen, along with the trends in each research category.

The increase in Government funding between FY 1983 and FY 1985 has been both in existing programs and in funding of new technology initiatives. The Strategic Computing Program, the Supercomputing Research Center, the Computer Measurement Research Facility, and the Energy Science Advanced Computation Programs are all new initiatives within the past two years and are funded at significant levels.

A wide range of jointly funded and coordinated projects to support research and development in high-performance computing is already in progress. A summary of the several major efforts follows.

1. Networking

Various network projects have been carried out jointly, or the results have been shared. The ARPANET/MILNET is used by DoD, DOE, NASA, NBS, NSF, and other parts of the Government. The CSNET activity of NSF is a joint effort supporting the entire computer science research community. The various networking activities continue to be closely coordinated as new networks become available.

2. Foundry Services

A joint effort has been undertaken by the National Science Foundation and the Defense Advanced Research Projects Agency that allows U.S. universities to use the DARPA-developed MOSIS fast turnaround VLSI implementation facility for university-based research and educational programs requiring the fabrication of digital designs as integrated circuit chips. Designs submitted by DARPA or NSF researchers in digital form over the ARPANET or Telenet using standard design rules and a standard artwork format are fabricated and returned to designers in four to six weeks. Both NMOS and CMOS are widely available, along with printed circuit board services. With appropriate authorization, it is also possible to provide access to VHSIC technology. The fabrication service is well documented, and other Government-sponsored programs are being encouraged to use the system.

3. Resources

CRAY machines at several locations are being made available for shared use. It is expected that the NSF supercomputing sites will be available for shared usage and quid pro quo arrangements should be explored for access to other systems. In addition, plans are in progress to provide access, including ARPANET and CSNET systems, to new numerical and symbolic processors as prototype machines become available from the many Federal projects in machine architecture and machine intelligence.

4. Architecture

1. The 128-node Butterfly multiprocessor built by Bolt, Beranek, and Newman (BBN) for DARPA was selected by the University of Rochester for experimental research funded by NSF.

2. The Pixel Planes Graphics projects at the University of North Carolina was jointly funded by NSF and DARPA.

3. The 4.2 BSD Unix software developed by Berkeley for DARPA has been distributed to most major U.S. universities for their research.

4. Data flow research at MIT and the Cosmic Cube at California Institute of Technology is jointly supported by DOE and DARPA.

5. Of the 38 architecture programs listed in Tables 1 and 2, 13 (or 34 percent) are funded by multiple agencies.

## E. PERSONNEL

Additional highly qualified research scientists and engineers will be needed to sustain the desired growth rate in high-performance computing. Growth in university faculty is essential to ensure an increasing pool of quality graduate students. Further, the number of qualified technical personnel for research in university, industry, and Government laboratories, as well as those with leadership capabilities to act as program managers within the Government, must be increased. The management and direction, as well as the actual research efforts, of basic research are set by persons who must be technically qualified to assess the accuracy and the quality of a proposed research program. This is particularly true of the program management personnel in the various Federal agencies involved in high-performance computing research. However, the number of new Ph.D.'s in computer science and information processing has stayed relatively constant since 1976. The increasing number of commercial successes in the field will create attractive opportunities to drain away both scientific personnel and funding from basic research areas which could lead to an overall decline in the rate of progress in the field and divert new graduates away from academic careers.

It is too soon to tell whether the number of Ph.D.'s will increase as a result of the increasing number of Bachelor's and Master's degrees being granted, since students from those classes who may be seeking Ph.D.'s have not yet completed their study. However, given earlier statistics, it would appear that without some method of encouragement, the number of

Ph.D.'s in computer science and related fields will not increase significantly in the future. If advanced research in Very High-Performance Computing is to continue and expand, these numbers must be increased. Without question, many of those persons with Bachelor's degrees and the ability to continue further in their education choose to go into industry instead. Indeed, a particularly disturbing trend is that some colleges and universities are considering limiting class size in undergraduate computer science and electrical engineering since the demands in these areas are stretching faculty resources and leaving less time for research. It is not possible to predict with accuracy the consequences of such actions on the role of development in high-performance computing.

Universities have increased their salaries substantially in the last several years to make the university environment more attractive. Starting salaries at many institutions are now comparable to those of industry; however, there is a problem in salary compression at the higher ranks in universities. Salaries in the public sector are lagging behind the levels in both academia and industry. Retaining high-quality personnel in universities and Government in the face of strong industrial incentives presents a particularly important challenge.

Industry lures not only newly graduated scientists and engineers, but also members of the existing, limited pool of technically qualified researchers in Very High-Performance Computing Research. The impetus for these people to move to the private sector to capitalize upon their research work can easily be seen in the rapidly increasing availability of venture capital to initiate new start-up companies.

The problem within the Government has been acknowledged to a limited extent. One attempt to provide an incentive for individuals to join or remain with the Government, at lower pay-scale levels, has been made by the use of a pay differential for critical engineering disciplines. Also annual increases have been given early in certain critical specialty areas. While recent improvements in Government pay for engineers have been beneficial, we note that computer scientists, broadly described, comprise an even scarcer category of skills. Whether a computer scientist has an engineering degree depends upon cultural factors at particular universities and not factors intrinsic to the discipline or the required education. The definition of those who qualify for engineering pay differentials should be broadened if the Government is to compete with industry in attracting qualified computer scientists.

Access to high-performance computing facilities by a wide range of scientists and engineers is considered an absolute necessity for areas in which major advances are now completely dependent on access to the most advanced scientific computational resources. Dedicated supercomputer systems exist at a small number of Government laboratories and research centers. Many of these systems are already saturated, moreover, they serve only a small portion of the nation's research community. NSF has estimated that the demand for access to supercomputing resources by the research community now exceeds the capacity by a factor of three. The overall issue of access and procurement of high-performance computing facilities is addressed more fully by the FCCSET Panel on Supercomputer Procurement and Access chaired by DOE.

F. FINDINGS AND RECOMMENDATIONS

This panel compiled a series of interrelated findings and recommendations on the issues that should be considered in order to ensure that the United States retains its strong national capability in advanced computing technologies. Other FCCSET panels have studied procurement and access strategies to drive the commercial supercomputer market and make supercomputer resources available to researchers and thus these issues are not repeated here.

The following recommendations take into account the importance of basic research as well as a number of significant engineering issues that must be considered immediately to speed existing technologies from the research laboratory into the mainstream of applications. The recommendations will assure that the United States continues to be the prime source of computing technologies in the decades ahead.

The following are the findings determined by the panel during the course of the discussions:

1. **Federal support for Advanced Computer Research and Very High-Performance Computing has grown significantly in recent years and a vigorous research and development program is developing.** The Federal program in Very High-Performance Computing is a collection of individual efforts by various departments and agencies. This effort is critical to provide the advanced computing capabilities that will be needed in the future. Despite the growth in the past two years in the level of Federal funding and the number of individual programs being conducted in this field, there are still significant areas that are not funded but appear to have promise for exploitation.

2. **Despite significant engineering expertise in multiprocessors, and computer technology in general, there is a lack of underlying theoretical understanding of Very High-Performance Computing.** Efforts need to be expanded in understanding such areas as multiprocessor architectures, languages, software systems, algorithms, problem decomposition techniques, etc., as well as in the applications that will be executed on new multiprocessor based computers. Equally large gains in performance can come from attention to software and algorithms as from hardware speed and the performance of specific architectures. The Federal investment in this critical area has been insufficient.

3. **The agencies which sponsor research in Advanced and Very High-Performance Computing necessarily give priority to their own needs and mission-specific goals (e.g., reliability, hardening, speed, and power).** Basic research proposals are typically formulated by individual researchers and presented to various agencies. Although a typical effort may be funded by a single agency, the possibility of funding by multiple agencies allows for a greater latitude in options to finance a complete, coherent program. The optimum result of multiple-agency investment is technology or concepts that are applicable to a broad spectrum of mission-specific goals.

4. **Technical program managers coordinate research on a case-by-case basis, and the method is effective.** Within the Government, program managers keep each other well informed about their individual programs and goals. Regular exchange of information occurs among the agencies of the Government concerned with the specific research area, so that local optimization of specific programs occurs.

5. **Multiprocessor based computer architectures appear to be so fundamental to success in Very High-Performance Computing that a thorough investigation of all promising architectures in this area is warranted.** A wide variety of promising architectures have been identified. It is particularly desirable to develop machines with the intrinsic capability to provide both numeric and symbolic processing.

6. **As computer research programs mature, access to advanced computing facilities, hardware, and knowledgeable scientists and engineers becomes critical to the success of the program.** Given the limited availability of these key resources in the areas of Advanced Computer Research and Very High-Performance Computing, the level of coordination required must be increased to ensure the optimum allocation of these resources to a number of individual programs.

7. **Although Federally sponsored research is intrinsically useful to the Government by definition, industry often lacks the economic motivation to capitalize on the results of this research.** As a result, domestic research may be better used outside the United States than within it.

8. **There are too few newly qualified research scientists and engineers in the United States to sustain the desired growth rate in Advanced and Very High-Performance Computing research through the end of the decade.** Over the past decade, the number of graduating Ph.D.'s in computer science and engineering, as well as in the overall information processing field, has remained relatively constant.

9. **To make the academic environment more attractive, universities have increased their salaries substantially in the past two years.** At many institutions, starting salaries are now comparable with those of industry. Within the Government, as well as in the university sector, the problem of salary compression at the higher ranks exists. Further, Government salaries have lagged behind those in the academic and industrial sectors.

10. **The ability to evaluate the performance of multiprocessor architectures is the key to sound, long-term research in this area, and may provide a solid foundation for their application.** Current performance evaluation and modeling techniques, as well as standardized benchmarks, are inadequate for the emerging multiprocessor architectures.

11. **The underlying support structure, or infrastructure, for research in Advanced and Very High-Performance computer research forms a key aspect of the overall program in this area.** The availability of research facilities, access to remote computer facilities, rapid prototyping of computer architectures and supporting electronics and mechanical structures, and the development of common design tools, performance benchmarks, interoperability standards, and communication protocols, etc., are essential for the timely development of advanced computer structures and information processing technologies.

12. **Although the coordination of basic research programs has been successfully conducted at the program manager level, it is difficult to quantify the existing formal, semiformal, and ad-hoc interagency coordination.** Coordination efforts are not clearly visible outside of the participating organizations. Visibility into the coordinating activities for this critical area of research must be enhanced at the interagency level. As a National response is formulated to the perceived threat to United States leadership in information processing technology, the ability to effectively present Government activities in Advanced Computer Research is essential.

To address the issues raised in the panel's findings, the following set of recommendations have been developed:

1. **MAINTAIN A VIGOROUS, COORDINATED RESEARCH PROGRAM.**

   The United States should maintain a coordinated program of research and development in Advanced Computer Research and Very High-Performance Computing, in particular, to ensure our national defense, to foster scientific excellence, and to enhance economic competitiveness. This program is growing and should continue to do so.

2. **INCREASE EMPHASIS ON UNDERSTANDING FUNDAMENTAL ISSUES IN PARALLEL PROCESSING.**

If industry and Government are to be able to fully apply new capabilities in Very High-Performance Computing, increased emphasis must be placed on achieving a better understanding of the fundamental issues in parallel processing such as algorithm development, problem decomposition techniques, languages, operating systems, and interprocess communication and synchronization mechanisms.

3. **PROMOTE RESEARCH ACTIVITIES THAT APPLY TO A BROAD CLASS OF PROBLEMS.**

The Government should adopt a policy that ensures that Federally supported research is as generic as possible, commensurate with meeting unique Government needs. Federal support in Very High-Performance Computing should be strengthened through the addition of generic research initiatives above and beyond those that are mission-specific, particularly in support of the research infrastructure.

4. **DO NOT OVER-COORDINATE BASIC RESEARCH.**

The Panel further recommends that no effort be undertaken at this time to force a more global optimization of the Federal basic research expenditures in Very High-Performance Computing because the field is in its infancy and moving ahead very rapidly. It is the Federal Government's role to fund long-term, high-risk basic research. Multiplicity of funding sources in basic research has been one of the important mechanisms to assure promising research is supported. This is particularly important in rapidly changing fields.

5. **EXPLORE A DIVERSE SET OF ARCHITECTURES.**

A Federal investment strategy must be followed which will allow exploration of the many promising diverse parallel processing architectures at both the exploratory and advanced development stages. A great deal will be learned by investigating a large variety of real problems and applications on these new architectures. Special effort should be undertaken to identify underlying structures that may be common to both scientific and symbolic processing.

6. **COORDINATE EXPLORATORY MACHINE ARCHITECTURE DEVELOPMENT EFFORTS.**

Programs to develop novel, exploratory machines should have increased coordination because significant resources are required and the leverage from interagency cooperation is greatest. This is the largest absolute growth area in Advanced Computer Research, having increased from $20 million in 1983 to an estimated $55.2 million in 1985.

7. **IMPROVE TECHNOLOGY TRANSFER MECHANISMS FROM FEDERALLY SPONSORED RESEARCH TO THE COMMERCIAL SECTOR.**

Technology transfer mechanisms need to be developed to aid both industry and Government in applying the results of research and development in Very High-Performance Computing. The FCCSET Panel should assess the difficulties and successes in applying this research to mission-specific areas and commercial applications.

8. **DEVELOP PROGRAMS DESIGNED TO AUGMENT THE NUMBER OF TRAINED RESEARCHERS.**

The Government should develop explicit programs aimed at training students in the area of Very High-Performance Computing, and computer science and engineering more generally. Consideration should be given to graduate-level programs, and feeder programs, with the goal of at least quadrupling the number of graduates with advanced degrees within a decade. Each agency should be encouraged to provide funds for grants to young researchers to increase the number of potential leaders in the field. The Presidential Young Investigator (PYI) awards might be used as a model for postgraduate-level candidates, but efforts to stimulate leadership at all levels of the educational process should be strongly encouraged.

9. **TAKE STEPS TO ENSURE THAT COMPENSATION IS ADEQUATE TO RETAIN QUALIFIED RESEARCHERS IN THE PUBLIC AND ACADEMIC SECTORS.**

Congressional support should be sought for a program that seeks to compensate highly qualified technical personnel who are attracted to Government service but would otherwise be unavailable because of existing salary differentials between the Government and private sectors.

10. **DEVELOP EFFECTIVE PERFORMANCE MEASUREMENT AND MODELING TECHNIQUES.**

A high priority should be established for increased research on computer performance metrics and modeling, and standards for benchmarking. The possibility of international cooperation in the development of common standards should be investigated.

11. **INVESTIGATE THE INFRASTRUCTURE REQUIREMENTS TO SUPPORT THE RESEARCH COMMUNITY.**

There are a number of services and facilities that are commonly required in advanced research and development programs. The Panel recommends that the infrastructure requirements be explored with a goal of developing a specific investment strategy in this area, with support from both the public and private sectors.

12. **MAINTAIN A VISIBLE INTERAGENCY COORDINATION EFFORT.**

We believe that interagency coordination of Federal Very High-Performance Computing Research and development activities is valuable and should be continued. This FCCSET Panel could continue to function for this purpose.

G. CONCLUDING REMARKS

The technological performance of the United States economy is fully dependent upon the industrial sector. The orientation of private research and development funds is critical to bring new technologies to the marketplace. The United States has been able to retain its leadership in high-performance computing and information processing due to the strong domestic industrial base. A major element in the Government's role is in creating the environment to foster the innovations necessary to maintain this industrial base.

Policies, when translated into specific programs, are the mechanisms to create the required environment, especially for basic and applied research, or exploratory development. Government policies fall into five areas influencing innovation (reference, "Federal Support for R&D Innovation, CBO Study, April 1984), such as:

1. **Microeconomics** - Fostering positive economic growth;

2. **Competition** - Minimizing national and international disincentives for innovation, such as trade barriers to protect U.S. industry;

3. **Tax Incentives** - Policies to allow recovery of R&D expenses;

4. **Regulatory Policies** - Influence of policies that may divert corporate funds away from R&D programs; and

5. **Institutional and Informational Support** - Enhancement of technology transfer from Federally sponsored research and development to the private sector.

Government policies in each of these areas can have either positive or negative effects upon innovation. Effective policies will allow the Government to leverage its portion of the research and development funding to meet its long-term mission-specific goals.

This Panel has brought together the primary Federal organizations funding basic research and exploratory development in Advanced Computer Research to determine the Federal activities in Advanced Computer Research and Very High-Performance Computing. This has required an in-depth understanding of the broad spectrum of technologies involved as well as knowledge of the mission-specific goals of the organizations which sponsor this research. In reviewing the investment profiles for the various agencies sponsoring work in this field, it is obvious that although the funding has been substantial and is increasing, there are still significant areas of research that are not being explored. The recommendations that are presented in this report form a basis for a sustained and vigorous program of research and development that is required to maintain the preeminent position of the United States in Advanced Computer Research and information processing technology.

# APPENDIX A

## SUMMARY OF FEDERAL INVESTMENT IN BASIC RESEARCH AND EXPLORATORY DEVELOPMENT IN ADVANCED COMPUTER AND VERY HIGH-PERFORMANCE COMPUTING RESEARCH

TABLE A-1

FEDERAL INVESTMENT IN BASIC RESEARCH
AND EXPLORATORY DEVELOPMENT FOR
ADVANCED COMPUTER RESEARCH
(~~BY ORGANIZATION~~/DIVISIONS)
(IN MILLIONS, CURRENT YEAR DOLLARS)

| | TOTAL PROGRAM | | | UNIVERSITY COMPONENT | | |
|---|---|---|---|---|---|---|
| | FY83 ACTUAL | FY84 ACTUAL | FY85 ESTIMATED | FY83 ACTUAL | FY84 ACTUAL | FY85 ESTIMATED |
| **DOD** | | | | | | |
| DARPA | 60.9 | 90.1 | 124.1 | 44.8 | 67.4 | 86.2 |
| ARMY | | | | | | |
| CECOM | 2.5 | 6.0 | 6.0 | 0.2 | 0.2 | 0.3 |
| ARO | 2.6 | 6.0 | 6.9 | 2.5 | 5.7 | 6.5 |
| TOTAL ARMY | 5.1 | 12.0 | 12.9 | 2.7 | 5.9 | 6.8 |
| AIR FORCE | | | | | | |
| AFOSR | 6.5 | 7.3 | 8.2 | 5.2 | 5.8 | 6.6 |
| RADC | 8.1 | 13.1 | 12.6 | 1.2 | 2.0 | 1.9 |
| AFAL | 1.4 | 1.7 | 1.8 | 0.4 | 0.5 | 0.5 |
| TOTAL AIR FORCE | 16.0 | 22.1 | 22.6 | 6.8 | 8.3 | 9.0 |
| NAVY | | | | | | |
| ONR | 6.6 | 7.2 | 7.6 | 6.3 | 6.8 | 7.2 |
| NAVELEX | 5.0 | 6.1 | 6.7 | 0.4 | 0.5 | 0.6 |
| NAVMAT | -0- | 3.0 | 9.2 | -0- | -0- | -0- |
| TOTAL NAVY | 11.6 | 16.3 | 23.5 | 6.7 | 7.3 | 7.8 |
| SUPERCOMPUTING RESEARCH CENTER | -0- | -0- | 12.0 | -0- | -0- | -0- |
| **TOTAL DOD** | 93.6 | 140.5 | 195.1 | 61.0 | 88.9 | 109.8 |
| **NSF** | 37.8 | 44.0 | 51.1 | 37.8 | 44.0 | 51.1 |
| **NASA** | 20.3 | 19.5 | 20.9 | 8.1 | 7.8 | 8.4 |
| **DOC** | 7.8 | 8.2 | 8.8 | -0- | -0- | -0- |
| **DOE** | 13.9 | 14.7 | 22.4 | 5.6 | 6.6 | 10.7 |
| **TOTAL** | 173.4 | 226.9 | 298.3 | 112.5 | 147.3 | 180.0 |

TABLE A-2

FEDERAL INVESTMENT IN BASIC RESEARCH
AND EXPLORATORY DEVELOPMENT FOR
VERY HIGH-PERFORMANCE COMPUTER RESEARCH
~~(BY ORGANIZATION/DIVISIONS)~~
(IN MILLIONS, CURRENT YEAR DOLLARS)

| | TOTAL PROGRAM | | | UNIVERSITY COMPONENT | | |
|---|---|---|---|---|---|---|
| | FY83 ACTUAL | FY84 ACTUAL | FY85 ESTIMATED | FY83 ACTUAL | FY84 ACTUAL | FY85 ESTIMATED |
| **DOD** | | | | | | |
| DARPA | 5.9 | 21.4 | 38.0 | 4.2 | 15.0 | 21.9 |
| ARMY: | | | | | | |
| CECOM | -0- | -0- | -0- | -0- | -0- | -0- |
| ARO | 1.2 | 1.8 | 2.2 | 1.1 | 1.7 | 2.1 |
| TOTAL ARMY: | 1.2 | 1.8 | 2.2 | 1.1 | 1.7 | 2.1 |
| AIR FORCE: | | | | | | |
| AFOSR | 2.5 | 2.8 | 3.3 | 2.0 | 2.2 | 2.6 |
| RADC | -0- | -0- | -0- | -0- | -0- | -0- |
| AFAL | 1.4 | 1.7 | 1.8 | 0.4 | 0.5 | 0.5 |
| TOTAL AIR FORCE: | 3.9 | 4.5 | 5.1 | 2.4 | 2.7 | 3.1 |
| NAVY: | | | | | | |
| ONR | 6.6 | 7.2 | 7.6 | 6.3 | 6.8 | 7.2 |
| NAVELEX | 2.1 | 0.8 | 1.8 | -0- | -0- | -0- |
| NAVMAT | -0- | -0- | -0- | -0- | -0- | -0- |
| TOTAL NAVY: | 8.7 | 8.0 | 9.4 | 6.3 | 6.8 | 7.2 |
| SUPERCOMPUTING RESEARCH CENTER | -0- | -0- | 12.0 | -0- | -0- | -0- |
| TOTAL DOD | 19.7 | 35.7 | 66.7 | 14.0 | 26.2 | 34.3 |
| **NSF** | 7.6 | 9.5 | 11.6 | 7.6 | 9.5 | 11.6 |
| **NASA** | 4.3 | 4.3 | 5.9 | 1.5 | 1.5 | 2.1 |
| **DOC** | 0.1 | 0.5 | 1.0 | -0- | -0- | -0- |
| **DOE** | 4.9 | 7.8 | 15.7 | 2.0 | 3.5 | 7.9 |
| **TOTAL** | 36.6 | 57.8 | 100.9 | 25.1 | 40.7 | 55.9 |

TABLE A-3

RESEARCH AREA FUNDING SUMMARY FOR ADVANCED COMPUTER AND VERY HIGH-PERFORMANCE COMPUTING RESEARCH

(IN MILLIONS, CURRENT YEAR DOLLARS)

| | ADVANCED COMPUTER RESEARCH | | | VERY HIGH-PERFORMANCE COMPUTING COMPONENT OF ADVANCED COMPUTING | | |
|---|---|---|---|---|---|---|
| RESEARCH AREA | FY83 ACTUAL | FY84 ACTUAL | FY85 ESTIMATED | FY83 ACTUAL | FY84 ACTUAL | FY85 ESTIMATED |
| Computational Mathematics | 15.8 | 16.8 | 17.9 | 4.7 | 7.8 | 10.7 |
| Machine Architecture | 20.0 | 28.8 | 55.2 | 16.8 | 24.4 | 54.8 |
| Machine Intelligence and Robotics | 28.8 | 42.0 | 61.6 | 3.4 | 4.8 | 6.7 |
| Distributed Computing and Software Systems | 24.7 | 38.9 | 49.0 | 4.9 | 6.2 | 7.3 |
| VLSI Design and Special Purpose Computing | 24.2 | 27.0 | 28.3 | 3.1 | 8.1 | 10.1 |
| Data Management | 10.0 | 8.6 | 10.2 | 1.1 | 0.9 | 2.4 |
| Theoretical Computer Science | 9.2 | 11.9 | 16.0 | 0.9 | 1.1 | 2.1 |
| Network and Research Facilities | 34.1 | 46.9 | 53.7 | 1.3 | 3.7 | 5.5 |
| Performance Evaluation and Modeling | 6.6 | 6.0 | 6.4 | 0.4 | 0.8 | 1.3 |
| TOTAL PROGRAM | 173.4 | 226.9 | 298.3 | 36.6 | 57.8 | 100.9 |
| VHPC PERCENT OF ACR: | | | | 21.1 | 25.5 | 33.8 |

APPENDIX B

ADVANCED COMPUTING RESEARCH AREA FUNDING
SUMMARY BY ORGANIZATION
FISCAL YEARS 1983, 1984, and 1985

TABLE B-1

ADVANCED COMPUTING RESEARCH AREA FUNDING SUMMARY BY ORGANIZATION

FY 1983 - ACTUAL
(CURRENT DOLLARS, IN MILLIONS)

| | COMPUTATIONAL MATHEMATICS | MACHINE ARCHITECTURE (H/W & S/W) | MACHINE INTELLIGENCE AND ROBOTICS | DISTRIBUTED COMPUTING AND SOFTWARE SYSTEMS | VLSI DESIGN AND SPECIAL PURPOSE COMPUTING | DATA MANAGEMENT | THEORETICAL COMPUTER SCIENCE | NETWORK AND RESEARCH FACILITIES | PERF EVAL AND MODEL | TOTAL FUNDING |
|---|---|---|---|---|---|---|---|---|---|---|
| DOD | | | | | | | | | | |
| DARPA | -0- | 2.8 | 13.5 | 11.5 | 15.8 | 1.5 | 1.8 | 14.0 | -0- | 60.9 |
| ARMY | | | | | | | | | | |
| CECOM | -0- | -0- | 1.0 | 0.5 | -0- | 1.0 | -0- | -0- | -0- | 2.5 |
| ARO | 0.4 | 0.5 | 0.7 | 0.3 | 0.4 | -0- | 0.3 | -0- | -0- | 2.6 |
| AIR FORCE | | | | | | | | | | |
| AFOSR | 2.1 | 0.8 | 1.7 | 0.7 | 0.2 | 0.2 | 0.6 | -0- | 0.2 | 6.5 |
| RADC | -0- | 1.4 | 1.2 | 5.1 | -0- | 0.2 | -0- | 0.2 | -0- | 8.1 |
| AFAL | -0- | 0.7 | -0- | -0- | -0- | 0.6 | -0- | 0.1 | -0- | 1.4 |
| NAVY | | | | | | | | | | |
| ONR | 1.4 | 1.0 | 2.2 | 1.0 | 1.0 | -0- | -0- | -0- | -0- | 6.6 |
| NAVELEX | -0- | 1.7 | 0.1 | 0.8 | 1.0 | 0.6 | 0.1 | 0.7 | -0- | 5.0 |
| TOTAL DOD | 3.9 | 8.9 | 20.4 | 19.9 | 18.4 | 4.1 | 2.8 | 15.0 | 0.2 | 93.6 |
| NSF | 0.8 | 4.3 | 2.6 | 2.8 | 2.2 | 1.6 | 6.3 | 12.2 | 5.0 | 37.8 |
| NASA | 2.4 | 4.1 | 3.7 | 0.5 | 3.5 | 3.6 | -0- | 1.1 | 1.4 | 20.3 |
| DOC | 1.2 | -0- | 1.5 | -0- | -0- | -0- | -0- | 5.0 | 0.1 | 7.8 |
| DOE | 7.5 | 2.7 | 0.6 | 1.5 | -0- | 0.7 | 0.1 | 0.8 | -0- | 13.9 |
| TOTAL | 15.8 | 20.0 | 28.8 | 24.7 | 24.2 | 10.0 | 9.2 | 34.1 | 6.6 | 173.4 |

TABLE B-2
ADVANCED COMPUTING RESEARCH AREA FUNDING SUMMARY BY ORGANIZATION

FY 1984 ACTUAL
(CURRENT DOLLARS, IN MILLIONS)

| | COMPUTATIONAL MATHEMATICS | MACHINE ARCHITECTURE (H/W & S/W) | MACHINE INTELLIGENCE AND ROBOTICS | DISTRIBUTED COMPUTING AND SOFTWARE SYSTEMS | VLSI DESIGN AND SPECIAL PURPOSE COMPUTING | DATA MANAGEMENT | THEORETICAL COMPUTER SCIENCE | NETWORK AND RESEARCH FACILITIES | PERF EVAL AND MODEL | TOTAL FUNDING |
|---|---|---|---|---|---|---|---|---|---|---|
| DOD | | | | | | | | | | |
| DARPA | -0- | 12.2 | 21.6 | 14.7 | 17.8 | 1.5 | 2.0 | 20.3 | -0- | 90.1 |
| ARMY | | | | | | | | | | |
| CECOM | -0- | -0- | 1.0 | 2.0 | -0- | 1.0 | -0- | 2.0 | -0- | 6.0 |
| ARO | 0.5 | 0.5 | 3.3 | 0.9 | 0.5 | -0- | 0.3 | -0- | -0- | 6.0 |
| AIR FORCE | | | | | | | | | | |
| AFOSR | 2.5 | 0.8 | 1.7 | 0.7 | 0.3 | 0.2 | 0.7 | -0- | 0.4 | 7.3 |
| RADC | -0- | 1.2 | 2.0 | 8.2 | -0- | 0.6 | -0- | 1.1 | -0- | 13.1 |
| AFAL | -0- | 0.9 | -0- | -0- | 0.3 | 0.4 | -0- | 0.1 | -0- | 1.7 |
| NAVY | | | | | | | | | | |
| ONR | 1.5 | 1.0 | 2.3 | 1.3 | 1.1 | -0- | -0- | -0- | -0- | 7.2 |
| NAVELEX | -0- | 1.5 | 0.6 | 0.9 | 1.2 | -0- | -0- | 1.9 | -0- | 6.1 |
| NAVMAT | -0- | -0- | -0- | 3.0 | -0- | -0- | -0- | -0- | -0- | 3.0 |
| TOTAL DOD | 4.5 | 18.1 | 32.5 | 31.7 | 21.2 | 3.7 | 3.0 | 25.4 | 0.4 | 140.5 |
| NSF | 0.9 | 3.9 | 3.9 | 3.7 | 2.3 | 2.3 | 8.8 | 14.4 | 3.8 | 44.0 |
| NASA | 2.2 | 3.9 | 3.6 | 2.0 | 3.5 | 1.9 | -0- | 1.2 | 1.2 | 19.5 |
| DOC | 1.2 | -0- | 1.5 | -0- | -0- | -0- | -0- | 5.0 | 0.5 | 8.2 |
| DOE | 8.0 | 2.9 | 0.5 | 1.5 | -0- | 0.7 | 0.1 | 0.9 | 0.1 | 14.7 |
| TOTAL | 16.8 | 28.8 | 42.0 | 38.9 | 27.0 | 8.6 | 11.9 | 46.9 | 6.0 | 226.9 |

TABLE B-3

ADVANCED COMPUTING RESEARCH AREA FUNDING SUMMARY BY ORGANIZATION

FY 1985 ESTIMATED
(CURRENT DOLLARS, IN MILLIONS)

| | COMPUTATIONAL MATHEMATICS | MACHINE ARCHITECTURE (H/W & S/W) | MACHINE INTELLIGENCE AND ROBOTICS | DISTRIBUTED COMPUTING AND SOFTWARE SYSTEMS | VLSI DESIGN AND SPECIAL PURPOSE COMPUTING | DATA MANAGEMENT | THEORETICAL COMPUTER SCIENCE | NETWORK AND RESEARCH FACILITIES | PERF EVAL AND MODEL | TOTAL FUNDING |
|---|---|---|---|---|---|---|---|---|---|---|
| DOD | | | | | | | | | | |
| DARPA | -0- | 20.4 | 34.1 | 17.8 | 19.6 | 2.4 | 3.1 | 26.7 | -0- | 124.1 |
| ARMY | | | | | | | | | | |
| CECOM | -0- | -0- | 1.5 | 3.0 | -0- | 1.0 | -0- | 0.5 | -0- | 6.0 |
| ARO | 0.5 | 0.6 | 3.6 | 1.0 | 0.6 | -0- | 0.6 | -0- | -0- | 6.9 |
| AIR FORCE | | | | | | | | | | |
| AFOSR | 3.2 | 0.6 | 2.1 | 0.7 | 0.3 | 0.4 | 0.7 | -0- | 0.2 | 8.2 |
| RADC | -0- | 1.3 | 2.6 | 7.9 | 0.1 | 0.4 | -0- | 0.3 | -0- | 12.6 |
| AFAL | -0- | 0.8 | -0- | -0- | 1.0 | -0- | -0- | -0- | -0- | 1.8 |
| NAVY | | | | | | | | | | |
| ONR | 1.6 | 1.2 | 2.4 | 1.3 | 1.1 | -0- | -0- | -0- | -0- | 7.6 |
| NAVELEX | -0- | 1.6 | 0.6 | 0.9 | 1.2 | 0.3 | -0- | 2.0 | 0.1 | 6.7 |
| NAVMAT | -0- | 0.4 | -0- | 8.8 | -0- | -0- | -0- | -0- | -0- | 9.2 |
| SUPERCOMPUTING RESEARCH CENTER | — | 12.0 | — | — | — | — | — | — | — | 12.0 |
| TOTAL DOD | 5.3 | 38.9 | 46.9 | 41.4 | 23.9 | 4.5 | 4.4 | 29.5 | 0.3 | 195.1 |
| NSF | 1.0 | 3.7 | 5.1 | 4.5 | 2.4 | 3.1 | 11.3 | 16.0 | 4.0 | 51.1 |
| NASA | 1.4 | 4.6 | 7.0 | 0.9 | 2.0 | 1.7 | -0- | 2.2 | 1.1 | 20.9 |
| DOC | 1.2 | 0.0 | 1.6 | -0- | -0- | -0- | -0- | 5.0 | 1.0 | 8.8 |
| DOE | 9.0 | 8.0 | 1.0 | 2.2 | -0- | 0.9 | 0.3 | 1.0 | -0- | 22.4 |
| TOTAL | 17.9 | 55.2 | 61.6 | 49.0 | 28.3 | 10.2 | 16.0 | 53.7 | 6.4 | 298.3 |

APPENDIX C

VERY HIGH-PERFORMANCE COMPUTING RESEARCH
AREA FUNDING SUMMARY BY ORGANIZATION
FISCAL YEARS 1983, 1984, and 1985

TABLE C-1

VERY HIGH-PERFORMANCE COMPUTING
RESEARCH AREA FUNDING SUMMARY BY ORGANIZATION

FY 1983 - ACTUAL
(CURRENT DOLLARS, IN MILLIONS)

| | COMPUTATIONAL MATHEMATICS | MACHINE ARCHITECTURE (H/W & S/W) | MACHINE INTELLIGENCE AND ROBOTICS | DISTRIBUTED COMPUTING AND SOFTWARE SYSTEMS | VLSI DESIGN AND SPECIAL PURPOSE COMPUTING | DATA MANAGEMENT | THEORETICAL COMPUTER SCIENCE | NETWORK AND RESEARCH FACILITIES | PERF EVAL AND MODEL | TOTAL FUNDING |
|---|---|---|---|---|---|---|---|---|---|---|
| DOD | | | | | | | | | | |
| DARPA | -0- | 2.8 | 0.3 | 1.1 | 1.2 | -0- | -0- | 0.5 | -0- | 5.9 |
| ARMY | | | | | | | | | | |
| CECOM | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | | -0- |
| ARO | 0.4 | 0.2 | 0.2 | 0.1 | 0.1 | -0- | 0.2 | -0- | -0- | 1.2 |
| AIR FORCE | | | | | | | | | | |
| AFOSR | 1.0 | 0.6 | 0.4 | 0.2 | 0.1 | 0.1 | 0.1 | -0- | -0- | 2.5 |
| RADC | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- |
| AFAL | -0- | 0.7 | -0- | -0- | -0- | 0.6 | -0- | 0.1 | -0- | 1.4 |
| NAVY | | | | | | | | | | |
| ONR | 1.4 | 1.0 | 2.2 | 1.0 | 1.0 | -0- | -0- | -0- | -0 | 6.6 |
| NAVELEX | -0- | 1.7 | -0- | 0.2 | -0- | 0.2 | -0- | -0- | -0- | 2.1 |
| TOTAL DOD | 2.8 | 7.0 | 3.1 | 2.6 | 2.4 | 0.9 | 0.3 | 0.6 | -0- | 19.7 |
| NSF | -0- | 3.4 | -0- | 2.0 | 0.7 | 0.1 | 0.6 | 0.7 | 0.1 | 7.6 |
| NASA | 0.4 | 3.7 | -0- | -0- | -0- | -0- | -0- | -0- | 0.2 | 4.3 |
| DOC | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | 0.1 | 0.1 |
| DOE | 1.5 | 2.7 | 0.3 | 0.3 | -0- | 0.1 | -0- | -0- | -0- | 4.9 |
| TOTAL | 4.7 | 16.8 | 3.4 | 4.9 | 3.1 | 1.1 | 0.9 | 1.3 | 0.4 | 36.6 |

TABLE C-2

VERY HIGH-PERFORMANCE COMPUTING
RESEARCH AREA FUNDING SUMMARY BY ORGANIZATION

FY 1984 ACTUAL
(CURRENT DOLLARS, IN MILLIONS)

| | COMPUTATIONAL MATHEMATICS | MACHINE ARCHITECTURE (H/W & S/W) | MACHINE INTELLIGENCE AND ROBOTICS | DISTRIBUTED COMPUTING AND SOFTWARE SYSTEMS | VLSI DESIGN AND SPECIAL PURPOSE COMPUTING | DATA MANAGEMENT | THEORETICAL COMPUTER SCIENCE | NETWORK AND RESEARCH FACILITIES | PERF EVAL AND MODEL | TOTAL FUNDING |
|---|---|---|---|---|---|---|---|---|---|---|
| DOD | | | | | | | | | | |
| DARPA | -0- | 12.2 | 1.3 | 1.2 | 4.8 | -0- | -0- | 1.9 | -0- | 21.4 |
| ARMY | | | | | | | | | | |
| CECOM | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- |
| ARO | 0.4 | 0.3 | 0.5 | 0.2 | 0.2 | -0- | 0.2 | -0- | -0- | 1.8 |
| AIR FORCE | | | | | | | | | | |
| AFOSR | 1.3 | 0.6 | 0.4 | 0.2 | 0.1 | 0.1 | 0.1 | -0- | -0- | 2.8 |
| RADC | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- |
| AFAL | -0- | 0.9 | -0- | -0- | 0.3 | 0.4 | -0- | 0.1 | -0- | 1.7 |
| NAVY | | | | | | | | | | |
| ONR | 1.5 | 1.0 | 2.3 | 1.3 | 1.1 | -0- | -0- | -0- | -0- | 7.2 |
| NAVELEX | -0- | 0.1 | -0- | 0.4 | 0.2 | 0.1 | -0- | -0- | -0- | 0.8 |
| NAVMAT | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- |
| TOTAL DOD | 3.2 | 15.1 | 4.5 | 3.3 | 6.7 | 0.6 | 0.3 | 2.0 | -0- | 35.7 |
| NSF | -0- | 2.9 | -0- | 2.5 | 1.4 | 0.1 | 0.8 | 1.7 | 0.1 | 9.5 |
| NASA | 0.6 | 3.5 | -0- | -0- | -0- | -0- | -0- | -0- | 0.2 | 4.3 |
| DOC | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | 0.5 | 0.5 |
| DOE | 4.0 | 2.9 | 0.3 | 0.4 | -0- | 0.2 | -0- | -0- | -0- | 7.8 |
| TOTAL | 7.8 | 24.4 | 4.8 | 6.2 | 8.1 | 0.9 | 1.1 | 3.7 | 0.8 | 57.8 |

TABLE C-3

VERY HIGH-PERFORMANCE COMPUTING
RESEARCH AREA FUNDING SUMMARY BY ORGANIZATION

FY 1985 ESTIMATED
(CURRENT DOLLARS, IN MILLIONS)

| | COMPUTATIONAL MATHEMATICS | MACHINE ARCHITECTURE (H/W & S/W) | MACHINE INTELLIGENCE AND ROBOTICS | DISTRIBUTED COMPUTING AND SOFTWARE SYSTEMS | VLSI DESIGN AND SPECIAL PURPOSE COMPUTING | DATA MANAGEMENT | THEORETICAL COMPUTER SCIENCE | NETWORK AND RESEARCH FACILITIES | PERF EVAL AND MODEL | TOTAL FUNDING |
|---|---|---|---|---|---|---|---|---|---|---|
| DOD | | | | | | | | | | |
| DARPA | -0- | 23.8 | 2.5 | 1.3 | 5.3 | 1.0 | 0.6 | 3.5 | -0- | 38.0 |
| ARMY | | | | | | | | | | |
| CECOM | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- |
| ARO | 0.4 | 0.4 | 0.6 | 0.3 | 0.2 | -0- | 0.3 | -0- | -0- | 2.2 |
| AIR FORCE | | | | | | | | | | |
| AFOSR | 1.6 | 0.5 | 0.5 | 0.2 | 0.2 | 0.1 | 0.2 | -0- | -0- | 3.3 |
| RADC | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- |
| AFAL | -0- | 0.8 | -0- | -0- | 1.0 | -0- | -0- | -0- | -0- | 1.8 |
| NAVY | | | | | | | | | | |
| ONR | 1.6 | 1.2 | 2.4 | 1.3 | 1.1 | -0- | -0- | -0- | -0- | 7.6 |
| NAVELEX | -0- | 0.4 | -0- | 0.6 | 0.6 | 0.2 | -0- | -0- | -0- | 1.8 |
| NAVMAT | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- |
| SUPERCOMPUTING RESEARCH CENTER | -0- | 12.0 | -0- | -0- | -0- | -0- | -0- | -0- | -0- | 12.0 |
| TOTAL DOD | 3.6 | 39.1 | 6.0 | 3.7 | 8.4 | 1.3 | 1.1 | 3.5 | -0- | 66.7 |
| NSF | -0- | 3.3 | 0.2 | 3.0 | 1.7 | 0.2 | 1.0 | 2.0 | 0.2 | 11.6 |
| NASA | 0.8 | 4.4 | -0- | -0- | -0- | 0.6 | -0- | -0- | 0.1 | 5.9 |
| DOC | -0- | -0- | -0- | -0- | -0- | -0- | -0- | -0- | 1.0 | 1.0 |
| DOE | 6.3 | 8.0 | 0.5 | 0.6 | -0- | 0.3 | -0- | -0- | -0- | 15.7 |
| TOTAL | 10.7 | 54.8 | 6.7 | 7.3 | 10.1 | 2.4 | 2.1 | 5.5 | 1.3 | 100.9 |

# APPENDIX D

# FEDERAL RESEARCH ACTIVITIES IN COARSE- AND FINE-GRAIN ARCHITECTURES

TABLE D-1

FEDERAL RESEARCH ACTIVITIES IN
COARSE-GRAIN ARCHITECTURES
(CHARACTERIZED BY INTERCONNECTION NETWORK)

| PROGRAM | SPONSOR | PERFORMER |
|---|---|---|
| PACKET SWITCHING | | |
| Ultracomputer | DOE/NSF | NYU |
| CEDAR | DOE/NSF | Illinois |
| FMP | NASA | Stanford |
| Homogeneous Machines | DARPA/NASA/DOE | Cal Tech |
| PUMPS | NSF | Purdue |
| Static Dataflow | NASA/DOE/NSF | MIT |
| Tagged Token Dataflow | DARPA | MIT |
| Ring Dataflow | NASA/DOE | Lawrence Livermore |
| CIRCUIT SWITCHING | | |
| TRAC | NSF/DOE/AFOSR | Texas |
| Butterfly | DARPA/NSF | BBN |
| TREE | | |
| DADO | DARPA | Columbia |
| AMPS | NSF | Utah |
| Reduction Tree | NSF | UNC |
| NEAREST NEIGHBOR | | |
| Finite Element Machine | NASA | NASA/Langley |
| Navier Stokes Machine | NASA | Princeton |
| Wavefront Array | NASA/ONR | USC |
| DAISY IV | NSF | USC |
| Cosmic Cube | DOE/DARPA | Cal Tech |
| CROSSBAR | | |
| Database Machine | DOE/NSF | Wisconsin |
| Multi-Micros | DOE | Los Alamos |
| S-1 | NAVELEX | Lawrence Livermore |
| RING | | |
| ZMOB | NSF/AFOSR | Univ. of Maryland |
| Psuedo-Ring | NSF | UC, Santa Barbara |
| Crystal | NSF | Wisconsin |
| MISC | | |
| PASM | ONR | Purdue |
| MD/C | ONR | Princeton |
| Parallel Speech | NSF | Purdue |
| Special Purpose Array Processor | NSF | Northwestern |
| Special Purpose Array Processor | NSF | Ohio State Univ. |
| Speech Architecture | NSF | Brown |

TABLE D-2

FEDERAL RESEARCH ACTIVITIES IN FINE-GRAIN ARCHITECTURES

| PROGRAM | SPONSOR | PERFORMER |
| --- | --- | --- |
| Connection Machine | DARPA | Thinking Machines |
| Boolean Vector Machine | NSF | Duke University |
| Massivelly Parallel Processor | NASA/Goddard | Goodyear |
| Non-Von | DARPA | Columbia |
| Blue Chip | ONR | Purdue University |
| Programmable Systolic Array | DARPA | Carnegie-Mellon University |
| Pixel Planes Processor | DARPA/NSF | University of North Carolina |
| Pipeline Bit-Serial Processor Array | NSF | Duke |

# APPENDIX E

# ADDITIONAL CONTRIBUTORS TO THE REPORT

## ADDITIONAL CONTRIBUTORS TO THE REPORT

A. UNITED STATES ARMY

Army Research Office
  Dr. Jimmie Suttle
CENTACS
  Dr. James Schell

B. UNITED STATES NAVY

Naval Electronics Systems Command
  Mr. John Machado
Office of Naval Research
  Dr. Paul Schneck

C. UNITED STATES AIR FORCE

AFOSR
  Dr. David Fox
  Major John Thomas
Avionics Laboratory
  Dr. Donald Moon
Rome Air Development Center
  Colonel David Carlstrom
  Mr. Sam DiWitte
  Mr. Richard Metzger
Systems Command
  Lietuenant Colonel James Riley

D. DEFENSE ADVANCED RESEARCH PROJECTS AGENCY
  Dr. Craig I. Fields

E. DEPARTMENT OF ENERGY
  Dr. James Decker

F. NATIONAL SCIENCE FOUNDATION
  Mr. Kent Curtis