DOCUMENT RESUME

ED 367 315 IR 016 595

TITLE	High Performance Computing and Communications Program. Hearing before the Subcommittee on Science of the Committee on Science, Space, and Technology. U.S. House of Representatives, One Hundred Third Congress, First Session (October 26, 1993).
INSTITUTION	Congress of the U.S., Washington, DC. House Committee on Science, Space and Technology.
REPORT NO	ISBN-016-043524-2
PUB DATE	94
NOTE	110p.
AVAILABLE FROM	U.S. Government Printing Office, Superintendent of Documents, Congressional Sales Office, Washington, DC 20402.
PUB TYPE	Legal/Legislative/Regulatory Materials (090)
EDRS PRICE	MF01/PC05 Plus Postage.
DESCRIPTORS	*Computer Networks; Computer System Design; Federal Legislation; Futures (of Society); Hearings; *Industry; Information Networks; Information Systems; *Information Technology; *Technological Advancement; *Telecommunications
IDENTIFIERS	Congress 103rd; High Performance Computing Act 1991; *High Performance Computing and Commun Program

ABSTRACT

This hearing explores how the High Performance Computing and Communications Program (HPCC) relates to the technology needs of industry. Testimony and prepared statements from the following witnesses on future effects of computing and networking technologies on their companies are included: (1) F. Brett Berlin, president, Brett Berlin Associates, accompanied by David R. Audley, managing director, Prudential Securities; (2) Peter R. Bridenbaugh, executive vice president, science, engineering, environment, safety and health, Aluminum Co. of America; (3) Paul E. Rubbert, unit chief, aerodynamics research, Boeing Commercial Airplane Group; (4) W. Donald Frazer, vice president, Massively Parallel Products, Oracle Corp; and (5) Marvin G. Bloomquist, manager, information technology, Mobil Exploration and Producing Technical Center. (JLB)

I F

HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM

ED 367 315

HEARING

BEFORE THE

SUBCOMMITTEE ON SCIENCE

OF THE

COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY U.S. HOUSE OF REPRESENTATIVES

ONE HUNDRED THIRD CONGRESS

FIRST SESSION

OCTOBER 26, 1993

[No. 78]

Printed for the use of the
Committee on Science, Space, and Technology

IR016595

U.S. GOVERNMENT PRINTING OFFICE
75–413 WASHINGTON : 1994

For sale by the U.S. Government Printing Office
Superintendent of Documents, Congressional Sales Office, Washington, DC 20402
ISBN 0-16-043524-2

2

COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY

GEORGE E, BROWN, Jr., California, *Chairman*

MARILYN LLOYD, Tennessee
DAN GLICKMAN, Kansas
HAROLD L. VOLKMER, Missouri
RALPH M. HALL, Texas
DAVE McCURDY, Oklahoma
TIM VALENTINE, North Carolina
ROBERT G. TORRICELLI, New Jersey
RICK BOUCHER, Virginia
JAMES A. TRAFICANT, Jr., Ohio
JAMES A. HAYES, Louisiana
JOHN S. TANNER, Tennessee
PETE GEREN, Texas
JIM BACCHUS, Florida
TIM ROEMER, Indiana
ROBERT E. (BUD) CRAMER, Jr., Alabama
DICK SWETT, New Hampshire
JAMES A. BARCIA, Michigan
HERBERT C. KLEIN, New Jersey
ERIC FINGERHUT, Ohio
PAUL McHALE, Pennsylvania
JANE HARMAN, California
DON JOHNSON, Georgia
SAM COPPERSMITH, Arizona
ANNA G. ESHOO, California
JAY INSLEE, Washington
EDDIE BERNICE JOHNSON, Texas
DAVID MINGE, Minnesota
NATHAN DEAL, Georgia
ROBERT C. SCOTT, Virginia
XAVIER BECERRA, California
PETER W. BARCA, Wisconsin

ROBERT S. WALKER, Pennsylvania*
F. JAMES SENSENBRENNER, Jr., Wisconsin
SHERWOOD L. BOEHLERT, New York
TOM LEWIS, Florida
HARRIS W. FAWELL, Illinois
CONSTANCE A. MORELLA, Maryland
DANA ROHRABACHER, California
STEVEN H. SCHIFF, New Mexico
JOE BARTON, Texas
DICK ZIMMER, New Jersey
SAM JOHNSON, Texas
KEN CALVERT, California
MARTIN R. HOKE, Ohio
NICK SMITH, Michigan
EDWARD R. ROYCE, California
ROD GRAMS, Minnesota
JOHN LINDER, Georgia
PETER BLUTE, Massachusetts
JENNIFER DUNN, Washington
BILL BAKER, California
ROSCOE G. BARTLETT, Maryland

ROBERT E. PALMER, *Chief of Staff*
MICHAEL RODEMEYER, *Chief Counsel*
KATHRYN HOLMES, *Administrator*
DAVID D. CLEMENT, *Republican Chief of Staff*

SUBCOMMITTEE ON SCIENCE

RICK BOUCHER, Virginia, *Chairman*

RALPH M. HALL, Texas
TIM VALENTINE, North Carolina
JAMES A. BARCIA, Michigan
DON JOHNSON, Georgia
ANNA G. ESHOO, California
EDDIE BERNICE JOHNSON, Texas
DAVID MINGE, Minnesota

SHERWOOD L. BOEHLERT, New York
JOE BARTON, Texas
SAM JOHNSON, Texas
NICK SMITH, Michigan
PETER BLUTE, Massachusetts

---

*Ranking Republican Member.

CONTENTS

WITNESSES

HIGH PERFORMANCE COMPUTING AND COMMUNICATIONS PROGRAM

TUESDAY, OCTOBER 26, 1993

HOUSE OF REPRESENTATIVES,
COMMITTEE ON SCIENCE, SPACE, AND TECHNOLOGY,
SUBCOMMITTEE ON SCIENCE,
Washington, DC.

The subcommittee met, pursuant to notice, in room 2318, Rayburn House Office Building, Hon. Rick Boucher (chairman of the subcommittee) presiding.

Mr. BOUCHER. The subcommittee will come to order.

This morning the Subcommittee on Science continues its oversight of the High Performance Computing and Communications Program as developed by the Federal Coordinating Council for Science, Engineering, and Technology in accordance with the High Performance Computing Act of 1991. The HPCC Program is a major component of the Administration's plan for creation of an advanced information infrastructure. The President's budget request for Fiscal Year 1994 was $1 billion for the program, an increase of 38 percent over the Fiscal Year 1993 funding level, indicating the importance the Administration attaches to this ongoing research initiative.

In previous hearings in March of 1992 and in February of this year, the subcommittee reviewed the management, operation, and future plans for the NSFNET from which the National Research and Education Network, mandated by the 1991 legislation, will evolve. In April the subcommittee reviewed the management of the overall HPCC Program.

Today we will explore how it relates to the technology needs of industry. The connection between the program and advancements in the computer and telecommunications industry is obvious, but the greatest long-term effect of the program will be in industries that make use of high performance computing and high speed networking. The potential value of these technologies to the Nation's economic strength was the main justification for the creation of the program and for its rapid growth.

Today we have asked witnesses from several industries other than the computer and telecommunications industries to describe how the computing and networking technologies are currently used by their companies and to project how these technologies will affect their companies in the future We've asked them to tell us whether research sponsored by the program is focused in areas that will be important for the future competitiveness of their companies and we solicit their views concerning the allocation of research funds among the major components of the program.

In addition, we would be pleased to receive from our witnesses recommendations for improvements in the implementation of the program. In particular, we're interested in the effectiveness of the National Coordination Office in communicating to the private sector actions with respect to the program and goals with respect to it and activities that are sponsored by the HPCC Program.

I am pleased to welcome our distinguished witnesses this morning, and we will turn to them after receiving statements from other members of the subcommittee. And we'll call first on the Ranking Republican Member, the gentleman from New York, Mr. Boehlert.

Mr. BOEHLERT. Thank you very much, Mr. Chairman.

The underlying issue at this morning's hearing was captured well in a recent cartoon by Jeff McNelly. It merely draws a gigantic cable labeled "information superhighway" strung along the telephone wires until it is connected to this somewhat ramshackle house. Inside a perplexed voice wonders, "Now how do we order pizza?"

I think the cartoon is right on target. The purpose of the HPCC Program is to come up with equipment and services that actually do assist people. We can't lose sight of that. This should be a program where the customer is always, or at least almost always, right.

This morning we'll be looking at the program's impact on industry, not individuals, and we'll be examining the entire HPCC Program, not just networking. But the same question needs to be asked: will all our research result in equipment and services that serve real needs?

I am not prejudging the answer to that. I've been a strong advocate of government involvement in high performance computing almost since the day I arrived in Congress more than a decade ago. To make such involvement work, we must do exactly what we're doing this morning: listen to the users of the program and make sure that federal agencies are doing the same.

I look forward to hearing from this morning's witnesses. I'm sure that by the end of this hearing we'll know not only how to use the information highway to order a pizza, but how to use it to make one.

And before concluding, I just want to note that Dr. Berlin is here and he has his special advisory team right behind him, two of his youngsters, including one of them—I want to alert all the panelists that Rebecca is taking notes.

So I look forward to this hearing and the expert witnesses. Thank you, Mr. Chairman.

Mr. BOUCHER. Thank you very much, Mr. Boehlert.

The gentleman from Minnesota, Mr. Minge.

Mr. MINGE. I have no comment.

Mr. BOUCHER. Thank you, Mr. Minge.

The gentlewoman from Texas, Ms. Johnson.

Ms. E.B. JOHNSON of Texas. Thank you, Mr. Chairman.

I simply want to express my interest and my appreciation for your having the hearing and express my interest in knowing more about the massively parallel processing and its practical uses in the industries other than the computing and telecommunications industry. So I will listen with great interest.

Thank you.

Mr. BOUCHER. Thank you very much.

We welcome now our distinguished panel of witnesses: Dr. David Audley, the Director and Manager of Strategic Analytics and Research for Prudential Securities in New York; Dr. Brett Berlin, president of Brett Berlin Associates of Alexandria, Virginia, and Chairman of the IEEE Committee on Computing and Applications Infrastructure; Dr. Peter Bridenbaugh, Executive Vice President, Science, Engineering, Environment, Health and Safety for the Aluminum Company of America; Dr. Paul Rubbert, Unit Chief, Aerodynamics Research of the Boeing Commercial Airplane Group from Seattle, Washington; Dr. W. Donald Frazer, Vice President, Massively Parallel Products of the Oracle Corporation from Redwood Shores, California; and Dr. Marvin Bloomquist, Manager for Information Technology of the Mobil Exploration and Producing Technical Center in Dallas, Texas.

Without objection, the prepared written statements of each of the witnesses will be made a part of the record. We would welcome your oral summaries of those statements, and in view of the number of witnesses we have, we would ask that you try to keep your statements to approximately five minutes. That will give us ample time to propound questions to you.

Dr. Berlin, if we may, we'd like to begin with you this morning, and we would welcome your statement.

STATEMENT OF DR. F. BRETT BERLIN, PRESIDENT, BRETT BERLIN ASSOCIATES, ALEXANDRIA, VA, ACCOMPANIED BY DR. DAVID R. AUDLEY, MANAGING DIRECTOR, PRUDENTIAL SECURITIES, NEW YORK, NY; DR. PETER R. BRIDENBAUGH, EXECUTIVE VICE PRESIDENT, SCIENCE, ENGINEERING, ENVIRONMENT, SAFETY AND HEALTH, ALUMINUM CO. OF AMERICA, PITTSBURGH, PA; DR. PAUL E. RUBBERT, UNIT CHIEF, AERODYNAMICS RESEARCH, BOEING COMMERCIAL AIRPLANE GROUP, SEATTLE, WA; DR. W. DONALD FRAZER, VICE PRESIDENT, MASSIVELY PARALLEL PRODUCTS, ORACLE CORP., REDWOOD SHORES, CA; AND DR. MARVIN G. BLOOMQUIST, MANAGER, INFORMATION TECHNOLOGY, MOBIL EXPLORATION AND PRODUCING TECHNICAL CENTER, DALLAS, TX

Dr. BERLIN. Thank you, Mr. Chairman, members of the committee. This is a particularly exciting time for us.

And, particularly, what I'd like to do is start by looking back 10 years. It's very appropriate that this hearing would be at this time; it's almost exactly 10 years since the first hearing that was ever held by this committee on supercomputing. It was actually held November 15, 1983, and that was significant because at that time even the scientific community was sharply divided on whether supercomputing was important.

There were some scientists who came and said that—such as Ken Wilson—that we would never be able to compete again if we didn't hurry up and do something about getting supercomputing out to the scientist. But there are others who said that it wasn't important, that minicomputers were fine, and that we just didn't need to move forward.

So as we move into this hearing, we'd like to take a look back at where we've come from and start out by talking a little bit about what the baseline is that we established at that time. There were four reasons why the high performance computing initiative in this committee decided to move.

First of all, U.S. leadership in the development and application of high performance computing was deemed fundamental and important to U.S. economic competitiveness and national security. And at that time we believed that science and technology was starting to fundamentally change the way it was happening, that we were moving from a experimentally-based society to a simulation-based society in science and technology. And what we see here today, Mr. Chairman, is evidence that not only that has happened, but, more importantly, that those changes have begun to infiltrate the ranks of database processing, management information systems, and have started to change the whole way we do decision support.

Now it's important to recognize that the changes we see today are not changes because the Cold War went away or anything else. It's because everything has gone to the—we have entered the information age in a whole new way and we will never go back. And at the end of this transition, as industry by industry begins the transition to high performance computing and modeling-based decisionmaking, they can never go back to the old ways. That's very important.

The second reason that this committee moved in 1983 was because they discovered that access to U.S.-made supercomputers could only be gotten by scientists if they left the United States. And, Mr. Chairman, you'll undoubtedly remember in Florida State University the rather dramatic testimony of Larry Smarr as he held up a "Scientific American" cover which had just been printed and said, "I'm proud to announce that I made this visualization. The only thing that I regret is I had to go to Germany to do it." And this committee recognized the importance of that and set in motion the National Science Foundation Center's program.

Third, this committee and the scientists and engineers who came before it recognized that there was a major competition brewing as other countries started to recognize what was happening. In the last 10 years that has—there have been dramatic changes in that. Toward the end of the eighties, for example, Japan had a rate of net new name customers that was several times the rate of new customers in the United States. They were all industrial customers. In 1985, they insisted to our Trade Ambassador that there were only four or five places in Japan that they could use supercomputers. By 1989, they had 140-some-odd supercomputers installed. That competition has now turned the corner and we are now starting to see the implications for that in competitiveness across the board.

And, finally, we discovered that the Government, based partly on its responsibility as a customer and its potential opportunity as a large sponsor of pre-competitive R&D, could have a major impact on America's overall competitive future.

I'd like to point out what has happened since then and particularly focusing on the NSF. I think we need to agree that the NSF, arguably, in their Center's program spent that money and it's some of the most profitable money the NSF spent during the 1980s. Each of the centers is well-run. It serves a wide variety of user groups, and I know, Mr. Chairman, you have met all the center directors and I think you'll agree that we would characterize them as a stubborn visionary or two, and that's been very important for the success of the centers.

It's also been highly leveraged as the private sector has gone to the centers, has used them, has changed the way they do business. We're seeing some examples of that here today, and you'll hear about that later.

The second aspect of that program was the modest investment in the INTERNET that the NSF began to make. Many people have forgotten that the purpose of the INTERNET was not just an experiment in communications; it was to connect supercomputers to each other and to the user base. And in the process of using those supercomputers we discovered a host of new applications, a host of new users, such that 40,000 networks are now connected to the INTERNET today. It's a tremendously successful and exciting program that this committee can take credit for, but also we shouldn't lose sight of.

I'd like to make just a brief comment about the three areas, Mr. Chairman, that you asked me to address; first of all, the relevance to industry needs. You're going to hear, the committee is going to hear testimony by my colleagues here about very specific ways that the HPCC is meeting industry needs. Ten years ago today, we had testimony that talked about applications that could use computers a thousand times of Cray-1. We are now almost at the point where we can see on the horizon that computer coming to fruition.

It's instructive to remember that those applications have not gone away. Not only have they not gone away, but we have created in the process, along the road, an entire new generation of applications, and it is my belief that by the end of the nineties we will have as many applications for terriflop computing that are fundamental to competitiveness as we have for gigaflop computing today.

As part of the role that high performance computing played in industry, we have to understand that there are a couple things, factors at work here. The first is that this is not just a trend of users who want to use more computing getting access to more computing. What we are seeing here is fundamental change in the way science, technology, decision support, and database management is done, and that's very important.

In the early eighties, SAAB as one example, was one of the first automobile manufacturers to use a Cray. They went from a VAX minicomputer directly to the use of a Cray because they had a fundamentally new application. The National Cancer Institute, when they put their supercomputer in, went directly from a VAX minicomputer to a Cray because there was a fundamentally different application.

As we look to the eighties and we—to the nineties and we see the database, large databases that are starting to aggregate and the need to optimize everything about the business environment, we are seeing the ground being laid for a whole new generation of these applications. And that, indeed, is what the nineties will be. It's an age of optimization. In the eighties we sought competitiveness in basic products. In the nineties everything that we compete on is optimized, and that includes not only airplanes, but it includes tennis shoes. And if you think about it now, tennis shoe advertisements emphasize high technology and optimized performance as opposed to just rubber from Malaya, and it's very important.

Finally, just the last couple of minutes I'd like to address briefly the role of the Federal Government in this, and I'd be happy to, in response to questions—I have a number of recommendations about the National Coordination Office, but in the interest of time I won't go into those.

But the federal role in the high performance computing comes, emanates from several different aspects. The first is that the Federal Government itself is a major user of the kind of things that high performance computing requires. There are many people who are walking or running around saying that, because defense needs have gotten less and the Department of Energy weapons programs have started to subside, that the federal requirement for high performance computing is getting less, but recall that we have everything from environmental cleanup to clean cars, to an entire plethora of research and development requirements that we can't even begin to address yet today with the computers we have. We also have a whole new generation of the need to optimize the way we do government and the way we do government services and databases, and those will require new generations of supercomputers.

And, finally, the government has several roles that are very important that we often underestimate. One which I'll just pull out of my testimony is the regulatory role. In order to meet the regulatory requirements of the nineties, the Federal Government is going to have to understand what it means to address optimized technologies. We're going to have to streamline and look at every one of the nuances of new products. Drug design is going to have to be regulated in whole new ways as we move to rational drug design on the civilian side, on the commercial side. So there are at least seven or eight reasons just from functional reasons why the Federal Government has a fundamental role and interest in high performance computing, and that's before we even address the issues of competitiveness.

Finally, I'll just close by pointing out that I think we're in one of those delightful times which I refer to as a 10-year itch. About every 10 years after a group of scientists or engineers has had to fight for large computing, they finally get a computer that in a small box equals what they fought for 10 years ago in a large box. And it happened in—in the fifties, you'll recall the famous report that A.D. Little put out that said there was only need for 15 computers in the world. In the sixties, it frustrated Thomas Watson to know that this little guy named Seymour Cray was building a computer out in the woods that was better than anything he could do, and why would anyone want one anyway? In the seventies it was minicomputers. Suddenly, minicomputers came about, and people who had fought to get large computers said: now I've got my departmental computer; I don't ever want to see a big computer again. That's when the Crays came on the scene and the large supercomputers, and people said: what are those for?

And then in the eighties we saw personal computers and people discovered they could do all sorts of things, and so why use a big computer? And by the nineties, now the issue is workstations, and workstations can do what a Cray-1 could do in 1976, but that's not the point. The point is that we—that our competitiveness does not depend upon what we could do in 1976; it depends upon what we can do in 1995.

Thank you.

[The prepared statement of Dr. Berlin follows:]

High Performance Computing and Communications: Checkpoints and Vision for 21st Century Competitiveness

by

F. Brett Berlin

Testimony before the U.S. House of Representatives, Committee on Science, Space and Technology Subcommittee on Science

October 26, 1993

Brett Berlin Associates

"There is nothing more difficult to plan, more doubtful of success, nor more dangerous to manage than the creation of a new systsm. For the initiator has the enmity of all who would profit by the preservation of the old system and merely lukewarm defenders in those who would gain by the new one."

Machiavelli, 1513

"Where there is no vision,
The people are unrestrained."

Proverbs (NASB)

Mr. Chairman, members of the Committee, thank you for inviting me to join in this hearing concerning the important role of high performance computing and communications in America's competitiveness as we engage the 21st century. As one who has been privileged to have been involved in both the industry and in the early development of the HPCC programs and initiatives, some of which actually started to take root and bloom as early as 1981, I am particularly pleased to be called on as part of the effort to assess where we have come and to set the course and vision for where we need -- or hope-- to go.

It is particularly appropriate that this "marker" hearing is being held at this time, only two weeks shy of the formal tenth anniversary of the first ever Congressional hearing on the subject of high performance computing. That hearing, convened by then Chairman Fuqua on November 15, 1983, established what has become a remarkable tradition of cooperation between the Congress, the President, industry and academia, now formally embodied in the HPCC Initiative and the HPCC Act of 1991. It also demonstrated the type of vision that has often been a hallmark of this Committee which has always been willing to square off with the complexities and subtleties inherent to incorporating technology into policy. Indeed, it was this Committee's willingness to put what even the majority in the science and technology community considered narrow and esoteric, into a long-range societal context that provided the spawning grounds for the development of the HPCC and Information Infrastructure vision that is now considered one of the hallmarks of former Committee member and now Vice President Al Gore.

I mention this anniversary because I believe it is important to measure our progress and to establish the roadmap for the next decade with a clear understanding of the baseline and the promises that motivated the original call to action.

<u>The Baseline</u>

The original motivations behind the eventual creation of the NSF HPC Centers, and the overall HPCC Initiative were based in four essential observations:

1. *U.S. leadership in development and application of HPCC technologies is important to U.S. and economic competitiveness and national security.*

2. *Access to U.S.-made supercomputers was out of reach of all but a select few within the entire technology R&D community.*

3. *There was a major competition brewing as other countries realized that HPCC was emerging as foundational to industrial competitiveness in the 1990's.*

4. *The government, particularly based on its responsibility as a customer and its potential opportunity as a large sponsor of precompetitive R&D, could have a major impact on America's overall competitive future in high performance computing, as a cooperative partner with the U.S. supercomputer industry.*

As a preamble to addressing the specific questions you posed in your invitation to testify, I wish to point out that there is much to celebrate when looking at what the program -- particularly that managed by the NSF -- has accomplished. In many respects, in fact, I think it is fair to argue that the money expended on the HPCC centers may be some of the most profitable investment NSF made during the 1980's. Each of the centers is well-run, serves a wide variety of user groups, and is led by a stubborn visionary or two. The NSF investment has been directly leveraged by both state and private funds, as well as considerable in kind and partnership support. Finally, all of the centers have made major contributions to both science and industrial competitiveness.

The modest investment in the INTERNET backbone, originally designed to provide access to the NSF HPCC Centers, has also been a major success for the nation. Since the backbone was put in place, over 40,000 networks worldwide have "plugged in" to the networked community. Within the U.S., where the network is growing fastest, the bulk of the new participants are entering via commercial gateways -- part of the vibrant new telecommunications and software industry and market that is now emerging. As only a small piece of the total INTERNET, the NSF backbone investment has to be one of the most highly leveraged of all Federal expenditures -- and certainly among the most effective.

On the other hand, there are many goals that have not been met. Some because of lack of focus, I suspect, but many more because we all discovered that success was much harder to come by. A few of the more difficult lessons include:

- The program has proven that moving to parallel computing is very hard. There is now also general agreement that the move is nonetheless essential.

- The program has demonstrated that designing new computers will still be hard for many years to come. Ten years ago, there were well over 100 proposed computer architectures; some predicted that the future would yield specially designed computers-on-a-chip optimized for each application. Building design tools that could easily create the complex masks required to prove that vision proved infeasible. Building systems (and designing software) based on standard microprocessors is not turning out to be much easier.

- The program has demonstrated that software and user priorities are still more important than architecture elegance; that because of the software issue delivered user performance rather than peak performance is the real figure of merit. There is now general agreement that software and applications need to be the focus for some time.

- Finally, the program has demonstrated the importance of regular, comprehensive, and objective oversight that is initiated from within the program. This program has many supporters across the country -- support that is generally well-deserved. However, even the program's strongest advocates were unable to defend it from a recent series of serious budget cuts by the Appropriations Committees. One reason is that there is no process for regular open review, definition and reassessment of objectives.

I now turn to the questions that the Chairman asked me to address for this hearing. As requested, I will keep my further remarks short so as to allow maximum time for questions and discussion.

Relevance to Industry Needs

High-end HPCC is demonstrably relevant to the needs of industry, as well as the continuing critical mission needs of government that industry contractors are being asked to address.
When this Committee first recommended that the NSF move forward to establish an HPCC program, U.S. national security, a number of other government mission functions, and a number of well-known commercial challenges were known to require computers up to 1000 times the power of the then state-of-the-art Cray-1. Even at the Cray-1 level of performance, supercomputers were considered a key element of the fundamental leading-edge tool set required for the maintenance of leadership in science, technology and leading-edge paradigm shifts in basic industries.

Dr. George Kozmetsky, one of the founders of Teledyne Corporation, and founder and director of the University of Texas Institute for Creative Capitalism, testified before this committee that:

> "The consequences of losing economic and scientific preeminence in the supercomputer industry are vast. The supercomputer is a central driver for the rapidly emerging worldwide computer/communications industry. It impacts communication developments, the renewal of basic industries, productivity increases and the development and expansion of new industries. It is essential in improving our educational structure, fulfilling critical manpower requirements and enhancing our industrial creativity and innovation. It is the seed for encouraging the emergence of a myriad of technology venture businesses in the context of a private enterprise system that has always been the unique American way to achieve and maintain U.S. economic and scientific preeminence."

In support of his conclusions, Kozmetsky identified twenty-one major areas, all key to competitiveness and national security, that were believed to be dependent on supercomputer technology development and leadership.

The NSF, in what is now referred to as the Lax Report, similarly identified a wide range of scientific and industrial applications which were already known to be constrained due to the lack of access to high performance computing or lack of mature software and algorithms essential to reliable results. Dr. Edward Knapp, then Director of the National Science Foundation, specifically highlighted a number of areas directly related to economic competitiveness and quality of life that were already starting to employ supercomputers. These included work in advanced materials, advanced electronics and circuitry, medical research, and macroeconomics.

In looking ahead it is important to note that we now have supercomputers in sight that will be capable of sustained performance equal to 1000 times the original Cray-1 -- the goal set by the community about ten years ago. During the intervening time, we have already seen major transformations in at least the following industries, specifically as a result of the U.S. high performance computing and communications industry: advanced research, aerospace, automobiles, oil, defense.

Additional industries that have begun, but have not progressed as far, generally due to the requirement for much higher capability, include: health care, chemical, pharmaceutical, environmental characterization and cleanup, advanced transportation modeling, law enforcement (identification), biotechnology, and communications. Finally, many experts anticipate that HPCC will be part of the total transformation of manufacturing process that could begin to take hold within the decade.

It is important to note that this is not simply a general list presented to give the impression that HPCC is a panacea. Technology itself does not solve problems, nor does its potential for application automatically mean it will succeed. In the case of HPCC, however, there are specific, economically and socially important applications that have been identified and in most cases demonstrated to be fundamental to the future economics and job sustainability of the industry. In health care, for example, approximately 25 cents of every dollar is spent on paper processing and information systems. There are a number of consortia currently working on various aspects of HPCC applications, particularly related to remote health care and consultation or similar activities. The University of Massachusetts, in partnership with several hospitals, a major health insurance company, medical records experts, and a leading supercomputer manufacturer, for example, is currently involved in a project designed to use high performance computing and leading edge natural language processing and retrieval techniques to automate some of the most expensive and error-prone aspects of the process of claims submittal and processing. In the pharmaceutical industry, Dr. Fred Houser, one of the leading developers of rational drug design approaches, now heads the first company in the industry specifically formed to develop drugs based on HPCC application in every part of the process. The competitiveness of Bionumeric is completely dependent on continued progress in the development and application of HPCC technology; every other pharmaceutical company is developing approaches and software that will allow them to fully integrate various, if not all of their development processes into an HPCC-based rational drug design process.

At this point, Mr. Chairman, before proceeding to the next section of this discussion, I think it would be useful to address the issue of the much-maligned "grand challenges" of science and technology, versus the currently politically correct "national challenge applications" proposed as more relevant to where we live (e.g., health care, distance and lifelong learning, etc.). There are many who have stated that the grand challenges are not sufficiently tied to the national economy and jobs to warrant such attention. While some of the original grand challenges were clearly in the realm of pure scientific endeavor (quantum chromodynamics, computational astrophysics, etc.), we should not lose sight of the fact that every one of the challenges identified in the FY93 program have substantial potential to contribute to economic competitiveness in at least one major industry or socially important program. One of the many significant shifts brought about by HPCC is the compression of the time between scientific discovery and economic viability. In other words, the days are gone when a scientist would have to wait years -- perhaps even a lifetime -- before seeing his work make a difference. Indeed, as HPCC becomes more ingrained into the scientific and industrial process, the potential for almost immediate insertion of new discovery into product designs could cause yet another revolution in the economic equation.

Appropriateness of the HPCC Program as a Federal R&D Activity

My colleagues across the country and I share the belief that the HPCC program, in general, represents an appropriate federal R&D activity. The simple fact is that HPCC technology and capability are fundamental to the quality of life -- underscored by economic and national security -- of 21st century industrialized nations. Economic competitiveness requires rapid turnaround of optimized products, global market understanding, and optimized business operations and decision making. These functions all increasingly require new generations of very high performance computing and communications to be available at appropriate price points. The increased complexity of national security -- both in terms of weapons and C3I (command, control, communications and intelligence) -- in the post-Cold War era also has heightened the importance of HPCC to the nation. With the Chairman's permission, I would like to submit for the record, after this hearing, the formal position statement developed by the IEEE underscoring the importance of the HPCC program as part of the ongoing Federal contribution to America's future technology leadership.

To many, the fact of the overarching importance of HPCC technology as a foundation for competitiveness is sufficient to warrant a major Federal role. Beyond the "clarion call" approach, however, there are a number of very pragmatic reasons that the Federal government now must be a major partner in the development and application of HPCC:

- First, the Federal Government remains a major current and potential customer. Challenges in operational weather and climate prediction, defense, energy, environment, health care, and a host of scientific research are dependent on emerging high performance computing capabilities. Furthermore, while parts of the traditional HPCC base, notably the nuclear weapons program, have diminished with the end of the Cold War, many of the new priorities will require even greater computing capabilities. New challenges in transportation, for example, could end up requiring extraordinary levels of computation and high performance communications. Environmental cleanup and health care information processing both are in their infancy in terms of computation.

- Second, most of the major Federal technology initiatives -- including efforts in advanced materials, global change research, advanced manufacturing -- as well as major government-industry collaborative research efforts -- such as those in aerospace and automobiles -- depend in some way on the cost-effective application of new generations of high performance computing. The recent White House announcement of the advanced automobile joint project with the USCAR consortium, for example, will focus on challenges that will easily require next generation HPCC systems. The automobile industry considers gains in HPCC so important, in fact, that they recently announced their own HPCC initiative to help ensure that the capability is in place to meet their future needs. Another area that is particularly dependent upon leading edge HPCC capabilities is medical research. Since the Federal government is a substantial player in this field, it is appropriate that the government participate directly with industry to develop the capability that this research requires.

- Third, money wisely spent in development of national HPCC capabilities can have leverage across a large percentage of national critical industries with a relatively small investment. HPCC methods and capabilities are rapidly causing an entire transformation of the global economy. Leadership that fosters aggressive application of HPCC is important to helping the economy bridge the cultural gap between old and new methods. A significant part of the current economic restructuring is due more to the accelerating transition to the information-based economy than to the shift of national priorities. We are moving from an experiment-based to a simulation-based economy; worker transition and retraining has to take this factor into account if we are to be successful in mitigating some of the transition pain.

- Fourth, the future of a substantial number of the Federal regulatory programs will depend on remote application of very complex models. Rational drug design and testing, for example, ultimately requires that the FDA be equipped to understand and properly assess advanced simulations as effectively as clinical trials. The DoD would like to be able to evaluate competing advanced weapons systems designs using new generations of precise, validated models. Aviation, consumer product safety, and environmental cleanup are other major arenas mitigating for timely involvement as a partner in industry's development of HPCC software and capabilities.

- Fifth, the Federal government shares with the private sector enormous challenges in dealing with very large databases that are growing exponentially. Not only are the sizes of these databases staggering by conventional terms, but the ancillary issues, such as security, privacy, etc., are concerns of all the users. While mass storage is usually thought of, when considering computation, as a scientific database problem, it is nonetheless true that the largest mass storage customers are commercial, non-scientific organizations. One of the nation's largest retailers is currently evaluating high performance computers as a strategy for better decision support. American Airlines recently installed a Kendall Square Research parallel processing system for an undisclosed major application. Neodata Corporation, the largest subscription fulfillment service, and EDS recently began an evaluation of a large scale parallel processing supercomputer to process complex transactions against a database scheduled to grow to approximately one trillion bytes! IRS, the Social Security Administration, and several other agencies face growing problems that eventually may require new generations of computational capabilities if they are to be effective.

- Sixth, HPCC holds out potential promises for development of innovative ways of upgrading the educational capabilities of the nation. The NSF centers have already demonstrated unique capabilities and concepts, and are involved in a number of experiments in this regard. I think it is fair to say that we are just scratching the surface thus far, in this area, and no one really knows where the investigation will ultimately lead.

Given the overall justification for involvement, the note of caution is that each activity and boundaries need to be carefully defined. Government does some things well, and is inherently impotent in other areas. Government, for example, can be an effective market stimulator by aggressive application of new technologies to solve real problems. It is a disaster when it tries to create an artificial market in hopes of "jump starting" a commercial company. While the principles are not sacrosanct, I would like to suggest that proper roles for government in HPCC will be characterized by four concepts:

- **Leadership**: The government must continue, as this Committee has done, to advocate vision, and to convene partnerships that will focus on addressing the problems to be solved. The NSF HPCC Centers program worked principally because the NSF provided leadership, but did not constrain the implementation. HPCC programs that focused on the methods rather than on the objectives have been less successful.

- **Regulatory Enablement**: The government must constantly seek to provide the regulatory environment conducive to the rapid growth of HPCC application within key industries. The issues here range from privacy, copyright, and the like; to tax and capital formation policy; to the impediments to procurement of leading-edge HPCC.

- **User Advocacy**: The government needs to be a strong advocate of its users, seeking to ensure that the taxpayers reap the economic results available by applying emerging computing capabilities in as many applications as possible. By doing this, the government also becomes an early market maker, as long as the users, not the research sponsors, are in the driver's seat.

- **Restraint**: The government can be a powerful partner. But the government can also inadvertently become a threat. The difference is often subtle, and can only be adequately discerned by careful consideration of industry perspectives at the front end of the planning process. But the difference is critical, often differentiating between "help" that is genuinely helpful, and that which can be lethal.

Appropriateness of The HPCC Technology Agenda

Concerning the appropriateness of the current HPCC program focus, I do not believe that I or anyone can make an adequate assessment. First, while the general directions of the funding are public, there is no document that delineates specifically how and to whom the funds are allocated. Second, everyone differs on what is needed because there has been no new baseline assessment. Therefore, I am reticent to specifically address the issue as you requested. However, I do believe there is a process we can define to get to a good answer.

When the HPCC interagency committee was first formed, under the auspices of OSTP using the FCCSET mechanism, its first task was to identify the areas of most pressing needs. It is important to recall that the committee was originally formed as a way of ensuring that the program reflected principally user priorities, rather than research dreams.

It is not surprising, therefore, that the first two areas of the funded program were designed to address the principal issue of providing a broad community of researchers with access to the most powerful capabilities available. To this end, the NSF formed its supercomputer centers and, with the help of DoE, NASA, and DoD, began to plan the INTERNET backbone. DoE and NASA also opened some of their systems capability to new users, and began an active collaboration with the NSF.

In addition to these initial efforts, the FCCSET committee set out to identify the key technology areas that were considered potential bottlenecks that could impede emerging applications of high performance computing. After considerable study, the committee agreed on the original research agenda and overall organization of the interagency effort; this agenda is basically the same one in effect today.

Meanwhile, on the industry side, virtually everything has changed in one way or another since the mid-1980's. A number of the companies that were considered major potential contenders, or at least exciting prospects, have gone out of business. Some architectural concepts that were highly touted have now been demonstrated to have limited utility; others now seem to have promise. None of the massively parallel processing vendors has turned the corner in terms of providing regular stable production science computing to a variety of customers. Only one of these vendors has recorded a profit, and even that vendor is still fragile. In terms of mass storage and semiconductors, the picture is much brighter than it was when the FCCSET committee made its first report.

In light of these changes, it seems that it is time for the HPCC to engage in a major reassessment, similar to that which was used to develop the original research agenda. A fresh roadmap, carefully taking into consideration a broad sample of industry and government users and industry technologists, could effectively ensure that the HPCC research allocations would be founded on as solid ground as they were when the program first came together, starting ten years ago.

Management of the HPCC Program and the Coordination Office

In most respects, the HPCC management has been a remarkable blend of agency programs and interagency cooperation rarely found over so long a period. Part of the reason for this success has been the essentially grass roots nature of the individuals involved and their longevity. Indeed, until last August, when the ARPA program manager was moved to another job within the agencies, each of the lead agencies had at least one committee member who has been in place for at least five years. In the case of ARPA, the same individual had been in place ten years. At DOE, the Chairman of the original FCCSET Committee, Dr. Jim Decker, is still in place at the Office of Energy Research, with his deputy, Dr. David Nelson, who has been on the committee ever since Decker stepped down. For over ten years, the only real problem that ever threatened the unity of the interagency group was a short period during which there seemed to be some disagreement at the cabinet level. Once this was resolved by White House intervention, the problem went away.

Despite the obvious collegiality within the FCCSET group, however, industry and the user community began to feel less and less able to provide constructive input. In essence, while every individual in the group would gladly meet with virtually any industry representative, many felt that issues were still not being resolved.

When the NCO was set up, consequently, many had expectations that it would act as a focal point for action on issues and problems. However, the NCO has neither power nor budget in an of itself, and is completely dependent on voluntary cooperation. When program issues are being considered, within the FCCSET process, it appears that this works. When the Director of the NCO tries to resolve an issue or get the agencies to do something they don't consider important, however, it is much more difficult.

While there is strong consensus that this way of running the operation is not optimal, there are also few options that would not violate the interagency realities that bureaucratically constrain the model in the first place. However, there are a few specific suggestions that I would offer:

- Require immediate appointment (within 30 days) of the HPCC Advisory Committee mandated two years ago. There is a serious need in the program for regular outside advice, review, and reporting. The Congress and the NCO both recognize this; the White House personnel system apparently does not. Some have suggested that one way of breaking the logjam would be to formally move the Board appointing authority to Dr. Gibbons. When HR 1757 goes to conference, perhaps the Chairman could consider this improvement.

- In addition to actively using the Advisory Board, the NCO needs to have a regular program of interaction with various segments, groups, constituencies. One possibility that a number of us have suggested is one or more workshops designed to solicit private sector detailed insights on a number of issues, ranging from the program mix to overall focus, to specific issues such as intellectual property. The NCO should be able to both approve and fund such workshops, rather than relying on the current consensus mechanism to get approval by committee.

- The NCO has recently begun disseminating its new program booklet via electronic mail, and has moved to include far more programmatic information. This effort should be commended and the NCO should be encouraged to act aggressively as a disseminator of information to any and all comers.

- The NCO should become the central clearing house for proposal and other information concerning HPCC-related solicitations. This could be a specific help to
- most companies. In particular, the NCO should be responsible for maintaining a single "bidder list" which would be automatically used by each of the participating agencies when announcing HPCCI-funded competitions. I understand that the NCO is moving to get the word out on BAA and related competitions. However, this should be a more formal charter.

- The NCO should be a central clearing house for a set of well-defined program goals and objectives applicable to each agency program. The NCO should be charged by the White House to engage a sub-panel of the Advisory Panel, or other Blue-Ribbon group, to review each program at least annually against its stated goals.

- Finally, the NCO should be empowered by the White House to report on all HPC expenditures and commitments, on a quarterly basis. This report should be publicly available.

In essence, the NCO should be chartered by the White House to look outward rather than inward. I believe that this is what Dr. Lindberg wants to do, and that he sees the need. The most glaring weakness of the program is that it is not open to scrutiny, and therefore is unable to take even constructive criticism well. The mission is too important, and the accomplishments are too good to allow such an Achilles heal to develop.

Conclusion: The Fundamental Choice and Vision

John Carlson, now the Chairman and CEO of Cray Research, underscored the relationship of supercomputing application to competitiveness in testimony presented to this Committee in 1984, put it this way:

> "Mr. Chairman, the world is undergoing a major revolution as supercomputing is applied across a broad spectrum of scientific and industrial applications. The only question at this juncture, is what nation will lead that revolution."

Today, as we reassess the vision and begin to implement course corrections, this warning still rings true.

Some will argue that HPCC has waned in importance. In fact, we are only at the beginning. And the next act of the play, the move towards the National Information Infrastructure, depends on our ability to keep the HPCC program aimed for success.

Likewise, some will argue that, despite all the rationale, government should not be involved. Factually, however, that question is moot: HPCC has already become so fundamental that government cannot escape involvement. Our only question is what else that role should entail.

Because of these realities, it is critical that we move with dispatch, but with care -- to ensure that the government response is crafted to maximize the strength of the uniquely American entrepreneurial spirit and process, while equally empowering the academic, laboratory and government members of the partnership to perform their most effective roles.

Thank you, Mr. Chairman. I look forward to answering any questions you might have.

BBA Brett Berlin Associates

F. Brett Berlin

Mr. Berlin is the President of Brett Berlin Associates, a small consulting group established in 1986 to provide unified information technology, policy, competitiveness and business development strategy advice to entrepreneurial company, institution and government executives; and Vice Chairman of the Board of the Institute for Clinical Information, Inc., a corporation recently formed by a group senior experts in the application of information technology in health care to develop innovative and competitive strategies for the development and use of computer-based patient records and related health care informatics. In addition, he serves as Senior Editor of *Technology Transfer Business*, the principal publication sponsored by the Association of Technology Business Councils and the only major publication devoted entirely to collaborative technology transfer, R&D commercialization and related competitiveness issues.

Mr. Berlin currently Chairs the Computing and Applications Infrastructure Subcommittee of the IEEE-USA Committee on Communications and Information (CCIP). He also serves on the Technology and Infrastructure Committee of the American Electronics Association, and is also a member of the Scientific Computation Division Advisory Panel of the National Center for Atmospheric Research (NCAR).

From 1991 to 1993, Mr. Berlin served as the Chairman of the HPCC Consortium, a unique cross-industry group sponsored by the American Electronics Association, including academic, research, industry and user groups interested in the HPCC Initiative and related application-oriented programs, and continues on its steering committee. From 1989 to 1991, he served as the Senior Policy Fellow for the Research Consortium, Inc. (RCI), an international consortium of HPC suppliers, integrators, and users. From 1983 to 1986, he served on the Department of Commerce Electronics Industry Sector Advisory Committee (ISAC-5). In 1987, he helped establish and served as one of the charter members of the SDIO National Test Bed Systems Engineering Panel. He has also served on the SDIO Technology Applications Advisory Panel (Electronics), the Electronic Industries Association Government Division Board of Directors, as Chairman of the National Coalition for Science and Technology, and as special advisor to the AEA National Information Infrastructure Executive Task Force.

4008 Ellicott Street • Alexandria, VA 22304
tel: 703-998-5888 • fax: 703-820-7181 • email: berlin@hpcc.gov

Since founding his consulting practice, Mr. Berlin has left active consulting on two occasions to take on specially created, strategy and marketing executive positions with a client company. Most recently, from August 1991 through December, 1992, he joined Kendall Square Research Corporation, as Vice President of Strategy and Government Affairs. KSR is the designer and producer of a new generation of general purpose, shared memory parallel high performance computer systems. During his KSR tenure, he served on the corporate management team assembled by the company's founders to establish and implement its product launch, government affairs and marketing strategies.

From October 1987 through October 1988, Mr. Berlin joined another client, Rockwell International, as Corporate Director, Strategic Program Development. In this position, he acted as a "one man think tank", chartered in the Washington Office to identify and analyze emerging issues key to the company's future technology (primarily defense) business, based on joint analysis of technology directions, emerging national/international political policy, global military strategy, and Rockwell defense and commercial corporate interests.

Prior to founding Brett Berlin Associates, Mr. Berlin was a Vice President and corporate officer of Cray Research, Inc., the world's leading manufacturer of general purpose supercomputers, where he served as the company's representative to official Washington and to the government R&D community, and as a principal strategist for government, prime contractor, and university marketing. During his tenure with Cray, he was the corporation's focal point for: science and technology policy, government procurement, trade, export control, and major supercomputing initiatives -- such as the National Science Foundation supercomputing centers program. Additionally, Mr. Berlin was senior corporate liaison to the Intelligence Community, and served as the Washington focal point for major new programs, such as the SDI.

From 1972 until the present, Mr. Berlin has been directly associated with the U.S. Air Force, on active duty until 1980, and as a reservist to the present. In this capacity, Mr. Berlin has served in a variety of technical, managerial, and consulting roles involving information and decision support systems, technical intelligence, analysis of performance of computer systems, advanced computer-based modeling, information systems research, and requirements planning for and acquisition of major information processing capabilities.

Mr. Berlin is a graduate of the U.S. Air Force Academy and The University of Texas Graduate School. Both graduate and undergraduate degrees are in Computer Science and Mathematics. His thesis, *Time-Extended Petri Nets*, focused on the performance evaluation of complex computer system architectures. His graduate advisor was Dr. J.C. Browne.

Mr. Berlin resides in Alexandria, Virginia, with his wife, Kathleen, and their three children, Rebecca, Christopher, and Stephanie.

4008 Ellicott Street • Alexandria, VA 22304
tel: 703-998-5888 • fax: 703-820-7181 • email: berlin@hpcc.gov

ENTITY POSITION STATEMENT

HIGH-PERFORMANCE COMPUTING AND COMMUNICATIONS

The Institute of Electrical and Electronics Engineers, Inc.
1828 L STREET, NW SUITE 1202, WASHINGTON, DC 20036-5104
(202) 785-0017

The High-Performance Computing and Communications (HPCC) program is essential to maintaining U.S. global leadership. We commend passage of the High-Performance Computing Act of 1991 and recommend rapid implementation of its elements.

Our future technology leadership relies upon these elements:

- Scientific research, development, and engineering are enhanced and transition time to market shortened when professionals have ready access to remote data sources. Such linkages enable complex data base analyses, couple human intellect to machine capabilities to optimize use of human pattern recognition capabilities, and help those professionals tackle problems that were not possible to solve without such high performance capabilities.

- Industrial design and manufacturing is more competitive and transition to a developed product is accelerated when using (1) simulations to evaluate paper designs more rapidly, accurately, exhaustively and at less cost; (2) visualizations and animations to enable insight to development and manufacturing challenges; and (3) Computer-Aided Design linked to Computer Integrated Manufacturing processes to eliminate barriers between engineering and manufacturing.

- High speed networks: (1) provide voice, video and data connectivity; (2) enable high-performance workstations with visualization and animation software to be linked to supercomputers; (3) interconnect computer mainframes to each other and to data storage peripherals within one data center; and (4) link supercomputers in different centers or even link networks of supercomputers to provide metacomputers for resolution of "Grand Challenge" problems.

These elements will help establish and enhance our competitive stance in the global economy, maintain economic viability and product excellence, and ensure the viability of critical national security systems.

IEEE United States Activities Board	July 1992
Approving Entity	Date

BACKGROUND

Traditional computer design strategies are running into basic physical limits, e.g. the so-called "von Neumann bottleneck." The machines cannot carry the burden of delivering trillions of computations per second. Computer designers are using new arrangements of computer elements as well as new technology to circumvent the limitations. Applied research and engineering needed to develop subsequent generations of computers is fraught with financial and technical risk.

U.S. firms at the leading edge of this technology tend to be relatively small -- sometimes too small to make the necessary investments in research without reasonable assurance of appropriate results. In order to attract substantial private investment, the high-performance computing market needs to be strong, predictable and based on well understood user requirements. The Federal government also needs to continue to increase its investment in research at this generic level. Furthermore, Federal funding at the generic applied research level must be focused, planned and dispersed in a manner that will stress the competitive posture of the American HPCC industry and user communities. In addition, by supporting local area networks of computers, Metropolitan Area Networks for small geographic regions, and wide area networks that link computer LANS, a meta machine can be assembled that is large enough to address the "Grand Challenge" issues once the distributed processing application software coordination problems are resolved. The high-performance computing community needs a long-term, well-funded and tightly coordinated Federal High-Performance Computing Program as part of the HPCC Act of 1991 to achieve its goals.

NREN will use transport media capable of handling a factor of ten to one hundred times the goal of 1 Gbps data rate. That data rate is sufficient to accommodate simultaneously 50 channels of broadcast quality HDTV, transmitting text or graphics on the network from two high-end performance workstations or from forty personal computers. A massive data volume would be generated, simulating air flow over a hypersonic aircraft's surfaces, if one assumes ten test points for each of a million grid points on the surface. Similarly, non-invasive radiosurgery treatment requires planning the positioning of radiation beams and the specification of their intensity such that a tumor receives 80% of the dosage and the surrounding tissues only harmless levels. In these cases, data may reside in remote hosts and require transmission to a local computer, marriage with local host data processing, and graphic portrayal for rapid assimilation by the user for decision-making. To avoid bottlenecks, network capacity must support such user demands.

Internet, a network of networks, already exists as an initial amalgam of Federal agency networks, private systems, state and regional networks and local research center and university networks. This pattern of networks will continue with NREN growing more complex as all the potential participants join the system. As the demands for connectivity and capacity by users grow, the challenges of leadership increase Learning while the system is relatively small enhances the probability of foreseeing the problems and structuring solutions before they become financially burdensome. In addition, alternate technical solutions can be assessed for segments of the network.

We recommend the following actions be considered in the HPCC program:

- Develop estimates of high-performance computing needs and available resources; refine theory and experience of how new computer architectures work and can be programmed and enhance computational methods and software operating systems, compilers and applications.

- Encourage the computing and communications industries and users to integrate supercomputing and parallel processing and high-speed data communications to make computer technology more affordable and accessible. Develop virtual networks of supercomputers and meta machines of networked computers. Focus, plan and disperse the technology investments in such a way as to stress a competitive posture for the American high-performance computing industry.

- Design the NREN (National Research and Education Network) to achieve data rates in excess of the 1 Gbps to accommodate projected user demand for data-rich applications, such as scientific visualization, and to support access to research devices, supercomputers and very large scientific and engineering data bases.

- Structure the NREN to serve as a leading-edge testbed for the development and study of basic technologies, applicable free space or guided network technology, high-level applications, standards, policies and network operational procedures using Internet, an amalgam of Federal agency networks, private systems, state and regional networks and local research center and university networks, as an initial testbed possibly leading towards a commercial network entity.

- Leverage Federal resources by requiring current and future HPCC participants to contribute some research, design, development, implementation or operational support to the expansion and use of the infrastructure created.

- Make early market insertion of new technologies, developed through both industry and government-supported resources, an HPCC goal. Focus on storage capability, data extraction and analysis software, optoelectric (e.g. amplifiers and star couplers) communication devices, and network management applications.

- Resolve legal issues regarding the protection of intellectual property, identification of liability under open access to systems, and personal privacy associated with service and product utilization data collection before the network is operational.

This statement was developed by the Committee of Communications and Information Policy of the United States Activities Board of The Institute of Electrical and Electronics Engineers, Inc. (IEEE), and represents the considered judgment of a group of U.S. IEEE members with expertise in the subject field. The IEEE United States Activities Board promotes the career and technology policy interests of the 250,000 electrical, electronics, and computer engineers who are U.S. members of the IEEE.

The leveraging of Federal resources occurs through joint Federal and private investments in research, development and product creation for the U.S. economy. One past example is a $12,000,000 DARPA investment that accelerated the creation of three high-performance workstation companies and the much larger Federal investment in the application of the technology. It can also result from the requirement that researchers, developers and users of NREN capabilities contribute to the process of creating a viable infrastructure, in addition to paying for use of the NREN capabilities.

Research and development as well as manufacturing consortia should be encouraged in technologies vital to the United States. Participants contribute researchers, facilities, administrative support, funds, and their dedication to achieving mutually beneficial products and services. The research is shared, but the product development remains on a competitive basis which results in the rapid insertion of the technology into the marketplace.

There are a variety of interfaces in the HPCC integration challenge and the required product enhancements, improvements or development. Large memory capacity, universal data handling tools, all-optical network components, the software to manage a hybrid network, and visualization workstations are key. HPCC helps the process of product development. Early market insertion assures availability of resources for development of future generations of new products.

Lastly, intellectual property, e.g., software licenses, can be misused by anonymous users on an open network. Similarly, access to host computers by network users, can be interpreted as introducing a user liability for misuse of the facility or for inducing harm to the user from system design/manufacturing errors. Hence, the liability cuts both ways. Permission to collect product and service usage data from a two-way transmission line to a facility as well as ownership and resale of that data, has system design and, of course, societal or privacy impacts. Operational requirements should be developed to provide guidance for system designers, developers, operators, providers and users.

Mr. BOUCHER. Thank you very much, Dr. Berlin.

Dr. Audley, we'd be pleased to hear from you. And if you would move the microphone over, that's good. Thank you.

Dr. AUDLEY. Mr. Chairman and members, I am David Audley, Director at Prudential Securities and Manager of the firm's Strategic Analytics and Research Department. I am pleased to be before you to discuss the high performance computing and communications program. Prudential Securities has been a beneficiary of some of the technologies that fall under the stewardship of the HPCC Program, and I'm prepared to report to you on our experiences.

In response to your objective to assess the relevance of the HPCC Program to the technology needs of U.S. industry, I will recount the manner and extent to which we at Prudential Securities use scalable high performance computing. In my prepared statement, submitted for the record, I have indicated some factual information about Prudential Securities that I'll admit at this point.

Prudential Securities was the first institution in the financial industry to make use of scalable high performance parallel, massively parallel computing. This occurred in 1988 with the introduction of an early generation of parallel supercomputer into our computing complex. This machine was a second-generation Intel IPSC/2, based on the 386 microprocessor.

This selection was made after an extensive evaluation of computing alternatives for Prudential's front office or strategic applications, an evaluation that included most alternatives from traditional supercomputers to networked workstations, and I'll say a little bit more about this evaluation process and how I think the general process we went through is so pertinent under the considerations of this committee.

Today Prudential's computing complex incorporates several of the most recent scalable high performance computers, as well as networked and clustered workstations. These machines are central to the firm's success and competitiveness. They are utilized 24 hours a day at more than 80 percent capacity by Prudential's activities around the world through commercially available communication links. These technical computers complement the mainframes, which are still the work horses for maintaining customer accounts, processing transactions, and supporting human resource applications.

What I'd like to do now is focus on a particular area that parallel computing and scalable high performance computing has made a difference to the competitiveness of our firm and how I think it has affected the country at large. This is in the area of mortgage and asset-backed securities. This is a major product area at Prudential Securities, and I will take a few moments to describe some of the essential elements of this market.

These markets seek to provide a safe investment vehicle for institutional investors while creating a ready and liquid market for collateralized debt. The largest sector of this market is the market of mortgage-backed securities. The mortgage market is the second largest investment market, second only to the market for U.S. Treasury securities. There are more than $2 trillion in debt outstanding in the mortgage market compared to the $3.45 trillion

75-413 O - 94 - 3

quoted for the U.S. Treasury market. This market is larger than the market for all corporate debt at $1.7 trillion.

In the mortgage market, pools of residential mortgages are used as investment vehicles for the Nation's public and private institutional investors: insurance companies, pension plans, and mutual funds, to name a few. This market is characterized by liquidity and efficiency, creating tremendous demand for mortgage collateral and thereby offering the American homeowner the cheapest source of mortgage lending in history.

In recent years the mortgage-backed securities market has seen tremendous innovation. Today institutions may obtain investments that have characteristics of cash flow that conform to their specific needs. This is particularly important for pension plans and insurance accounts that have defined future liabilities that have unique demands. The mortgage-backed securities market offers securities that are constructed to conform to the needs of these institutions with safety and return previously unavailable.

In the prepared statement I make a few remarks about how this market works. I think the most important thing that I should note at this time is that the growth in this market from the early eighties has generally paralleled the availability of analytics and fast database response to the securities industry, primarily the traders and bankers that support this market.

The growth of the mortgage-backed securities market has paralleled the cost performance progress in computers in the last 15 years. The American homeowner benefits directly from the competitiveness of this market and the resulting narrowing of the cost of mortgage financing from nearly 3 percent over comparable U.S. Treasury yields in 1986 to an average of only 1.6 percent over U.S. Treasury yields in the 1990s. So this is an improvement in bond parlance of 150 basis points or 140 basis points. That's 1.4 percent improvement in the efficiency of this market, which has generally been passed on to the American homeowner mortgage borrower.

What this means in terms of dollars and cents is that, for the average mortgage originated in the Washington metropolitan area, this is a return to the homeowner of more than $2,000 a year in taxable consumer available, spendable income. And this is not a reduction in mortgage rates that corresponds to the tremendous rally that we've seen in the treasury market. This is purely what we like to see in the increase of an industry's efficiency and effectiveness. And this is largely attributable since the mid-1980s to the availability of high performance computing to be able to match up mortgage borrowing pools and investors in institutional accounts.

The technology needs of U.S. industry is driven by competition—in the case of the competition for the—in this case, for the competition for the world's capital resources. The U.S. home mortgage industry is attracting capital flows from around the world. It is a safe, flexible, and efficient market for investment. This is increasingly made possible by high performance computing. Overseas institutions have come to Wall Street to see the expanded—the expanding use of technology in the capital markets and see tremendous barriers to competitive entry into American computer know-how.

Beyond home mortgages, the same financial technology has led to the securitization of automobile loans, leases, and consumer credit card debt, making capital available to American consumers and enterprises with unprecedented efficiency. The next arena for this technology is the securitization of health care receivables. The investment potential here can remove the assumption that taxpayers must pay the Nation's health care overhead. Instead, that overhead can capture the efficiencies of the capital markets in exactly the same way as has occurred in the mortgage market.

In assessing the technologies that have emerged over the years and are currently under the sponsorship of the HPCC Program, I can only say that as a professional who has responsibility for technology at a major commercial institution, there are three things that are important to me when I look at incorporating new technology into our business goal. The first of these, of course, is performance, but, moreover, the performance cost quotient not only today, but how that performance cost quotient goes out into the future. When we brought in our first parallel machine in 1988, we could see that there was an opportunity for scalability of these machines by both making them larger, but also the suggested scalability in technology. In fact, what we found was that as the 386 processor was replaced by the 860 processors, it was an immediate swap in our computer chassis to put in the 860 processors and to get the two and a half times performance increment using the same software that ran—or the source code that ran on the previous machine simply with a recompile. And we've seen that with each successive generation of technology that has been introduced into the parallel environment.

Mr. BOUCHER. Dr. Audley, we really do need to move along a little bit.

Dr. AUDLEY. Okay.

Mr. BOUCHER. I'm going to come back to you for some questions——

Dr. AUDLEY. All right.

Mr. BOUCHER. We are trying to keep these opening statements to five minutes, if we can.

Dr. AUDLEY. Very good, sir.

[The prepared statement of Dr. Audley follows:]

Testimony on High Performance Computing and Communications
Before the House Committee on Science, Space, and Technology
Subcommittee on Science
by David R. Audley, Ph.D.
Director, Prudential Securities Incorporated
October 25, 1993

Mr. Chairman and Members of the Subcommittee:

I am David Audley, Director at Prudential Securities and Manager of the firm's Strategic Analytics and Research Department. I am pleased to be before you to discuss the High Performance Computing and Communications (HPCC) program. Prudential Securities has been a beneficiary of some of the technologies that fall under the stewardship of the HPCC program and I am prepared to report to you on our experiences.

In response to your objective to assess the relevance of the HPCC program to the technology needs of U. S. industry, I will recount the manner and extent to which we at Prudential Securities use scalable, high performance computing. First, let me describe Prudential Securities and our business.

Prudential Securities Incorporated is a full-service, world-wide broker-dealer and investment bank. As a wholly owned subsidiary of the Prudential Life Insurance Company of America, it provides individuals, institutions, corporations and governments with such services as investment advice, asset management, securities brokerage, investment banking and retirement planning. It is the third largest full-service brokerage firm in the United States with more than 300 domestic offices, 17 in Europe, and 7 in the Pacific Rim.

A PIONEER IN SCALABLE, HIGH PERFORMANCE COMPUTING

Prudential Securities (PSI) was the first institution in the financial industry to make use of scalable, high performance (parallel) computing. This occurred in 1988 with the introduction of an early generation of parallel supercomputer into the PSI computing complex, the Intel iPSC/2. This selection was made after an extensive evaluation of computing alternatives for PSI's "front office", strategic applications -- an evaluation that included most alternatives from traditional supercomputers to networked workstations.

Today, the PSI computing complex incorporates several, scalable high performance computers as well as networked and clustered workstations. These machines are central to the firm's success and competitiveness. They are utilized 24 hours a day at more than 80% capacity by PSI activities around the world through commercially available communication links. These "technical" computers complement the mainframes which are still the workhorses for maintaining customer accounts, processing transactions, and supporting human resource applications.

The technologies of the High Performance Computing and Communications Initiative (HPCC) are redefining many elements of the capital markets. Prudential Securities is no longer alone in taking advantage of HPCC related technology -- some of our competitors have added scalable, high performance computers in the recent past. The following, however, describes one aspect of the technology-driven marketplace at Prudential Securities.

THE MORTGAGE AND ASSET BACKED SECURITIES MARKETS

Scalable, high performance computing is the cornerstone of PSI's competitive strategy in the capital markets. The firm has made a commitment to become a leader in the mortgage and asset-backed capital markets. These markets seek to provide safe investment vehicles for institutional investors while creating a ready and liquid market for collateralized debt. The largest sector of this market is the market of mortgage backed securities (MBSs). The mortgage market is the second largest investment market ($2.13 trillion), second only to the market for U. S. Treasury securities ($3.45 trillion) and larger than the market for corporate debt ($1.70 trillion). In the mortgage market, pools of (mainly) residential mortgages are used as investment vehicles for the nation's public and private institutional investors--insurance companies, pension plans, and mutual funds, to name a few. This market is characterized by liquidity and efficiency, creating tremendous demand for mortgage collateral, and thereby offering the American homeowner the cheapest source of mortgage lending in history.

In recent years, the mortgage backed securities market has seen tremendous innovation. Today, institutions may obtain investments that have characteristics of cash flow that conform to their specific needs. This is particularly important for pension plans and insurance accounts that have defined, future liability demands that are unique. The mortgage backed securities market offers securities that are constructed to conform to the needs of these institutions, with safety and return previously unavailable. Collateralized Mortgage Obligations (CMOs) are the securities created by the MBS market in response to institutional need. Essentially, these securities tailor the cash flow characteristics of the underlying mortgage collateral into manufactured "tranches" and transfer the remaining cash flows to other primary and companion classes of securities.

The tranched classes of security in CMOs are the consummate product in the spirit of the service principle of creating a tailored product in a batch of one. At the same time, the whole of the underlying collateral must be factored into tranches leaving little if any residual remainder. This process of security design and manufacture is done by the underwriter while at the mercy of the securities markets. The need for speed is critical for controlling risk and providing maximal value for investors. The millions of calculations needed to successfully complete such a transaction demand the highest levels of performance available in today's computers.

In addition to the primary market activity of creating new securities, the robustness of this market has been built on an active and liquid secondary market. Here, dealers must provide a bid for securities at competitive levels. As many of these securities are one-of-a-kind, databases and analytic tools are maintained to provide immediacy in the market for pricing, risk analysis and hedging. Without a ready secondary market for CMOs, this product could not exist. Again, the computational needs of data retrieval and analysis demand the highest levels of computational performance available today. These are real-time applications where speed is essential to compete successfully.

The growth of the MBS and CMO markets have paralleled the cost performance progress in computers in the last 15 years. The American homeowner benefits directly from the competitiveness of this market and the resulting narrowing of the cost of mortgage financing -- from nearly 3% over comparable U. S. Treasury yields in 1986 to an average of only 1.60% over U. S. Treasury yields in the 1990s.

The technology needs of US industry is driven by competition. In this case the competition is for the world's capital resources. The US home mortgage industry is attracting capital flow from around the world. It is a safe, flexible and efficient market for investment. This is increasingly made possible by high performance computing. Overseas institutions come to Wall Street to see the expanding use of technology in the capital markets and see tremendous barriers to competitive entry due to American computer know-how.

Beyond home mortgages the same financial technology has led to the securitization of automobile loans, leases, and consumer credit card debt making capital available to American consumers and enterprises with unprecedented efficiency. The next arena for this technology is the securitization of health care receivables. The investment potential here can remove the assumption that taxpayers must pay the nations health care overhead. Instead that overhead can capture the efficiencies of the capital markets in exactly the same way as has occurred in the mortgage market.

ASSESSING THE HPCC PROGRAM

In framing the following assessment, and in light of the above. the following claims are immediately rejected. First, that there are no successful commercial applications of scalable, high performance supercomputers. Second, these computers are not effective and are too hard to use. In fact, this technology has been found to be invaluable in the essential economic arena of the capital markets. It has been effective in increasing market efficiency and in reducing transaction costs. In addition, it was not too hard to take advantage of this technology. Prudential Securities has no special secrets or mystical insights. The firm's application developers seized upon the technology opportunity with the intent of capturing an advantage over our competitors.

At this point, an assessment of the HPCC initiatives would recognize the following as first and foremost.

> The Federal investment in HPCC technologies has had a positive impact on at least one commercial endeavor in an entirely unexpected way.

These technologies have and will continue to positively impact capital market efficiency and costs. It can be argued that of all the promised benefits for the technologies developed under the HPCC initiatives, the last one envisioned would have been in the economic area of home mortgage financing. This underscores the assertion, much maligned, that the biggest beneficiary of any technological advance could be entirely unforseen.

The HPCC initiative is relevant to the technology needs of American industry. The financial industry is but one case in point. Others can surely be similarly described. As the initiative has been relevant to the financial industry, it has also been important to Prudential Securities in advancing our competitive success in the marketplace. Beyond these observations, however, it must be recognized that the HPCC initiative is paramount as a federal R&D activity. It has brought about a new paradigm for computing. As to whether this will follow through to a true paradigm shift for general purpose computing, it may be too early to know. The potential has been demonstrated. The follow through must be completed. The essentials of the follow through involve general purpose applications and the accessibility of the new paradigm in computer research and education. The steps needed here are as follows.

> Involve independent software vendors in developing scalable applications.

Business systems may well benefit from scalable, high performance technologies in the area of database and human resource applications. This is the occurrence that could redefine the nature of the mainframe computer.

As an aside, it is notable that Oracle Corporation has methodically marched through the array of scalable, high performance computers that are commercially available and will introduce a version of its relational database management software (RDBMS) which will operate on virtually all major platforms. Independent software vendors must be more broadly involved in the new HPCC technologies. Hardware manufacturers, the independent software vendors, and the research establishment should be encouraged to seek alliances with commercial enterprises to further explore the promise of HPCC technologies in business settings.

Promote software application and interface standards.

There is much to be done here to encourage integrated environments that can utilize different architectures and strategies in supporting applications. Operating system disparity is but one issue. Just obtaining a usable common source, from what are thought of as portable languages, to run on more than one machine can be a challenge. This challenge is currently being met by the American National Standards Institute (ANSI) through the standardization of programming languages such as C. These efforts, however, are ad hoc and without industry-wide coordination. The HPCC initiative can do much to move these efforts along.

Provide broad access to scalable systems by the national research and education infrastructure.

This last item cannot be overstated. It is crucial to the nations competitiveness that graduates have experience in the new computing paradigm so that as employers move to capture the advantage of HPCC technologies, a talent pool of trained professionals will be available. This access invokes the need for the National Information Infrastructure (NII) program. This capability can bring the national computing resource to the research and education establishment. While scalable computers need to be widely available, the truly challenging problems of large scale computing may be best met by resources accessible through NII.

CONCLUDING COMMENT

The question is posed as to whether the HPCC initiative receives funding priority in the President's budget consistent with the importance of the program. Ours is a knowledge culture. Information processing, numerical computation and machine reasoning are the means of a new way of thinking. The HPCC initiative creates the resource to conduct the new thought processes of our age. How can we possibly afford to under-fund these initiatives when they may well provide the means to marshal the national intellect in bettering the human condition, be it through newly engineered drugs, an abundance of energy, or, indeed, housing accessible to all.

David R. Audley, Ph. D.
Director, Prudential Securities
Manager of Strategic Analytics and Research

David R. Audley is Director, Prudential Securities, and Manager of Strategic Analytics and Research. His department is responsible for the firm's analytic trading and sales systems, risk management, securities and market research, financial technologies, and a proprietary arbitrage trading account.

Dr Audley has been with Prudential Securities since 1987. Prior to that he was on active duty in the Air Force for 16 years. His last assignment was with the Strategic Defense Initiative Organization at OSD where he was Deputy Director for Battle Management, Command, Control, and Communications.

He is a 1968 graduate of The Citadel, holds a Master of Electrical Engineering Degree from the University of Southern California (1969), and received the Doctor of Philosophy degree from Johns Hopkins University in 1972.

Mr. BOUCHER. Dr. Bridenbaugh, we'll be pleased to hear from you.

Dr. BRIDENBAUGH. Thank you very much, Mr. Chairman and members of the subcommittee. I'd like to summarize three of the major points in my written statement, and I'd like to start with NSF Supercomputing Centers, which I believe have played and are playing a critical role in technology transfer between industry, universities, and the national labs. These centers give companies access to supercomputing capabilities that they wouldn't necessarily have on their own, and so we can now work on problems which are either too large for the in-house computers or are time-critical and have to be solved very rapidly. And, more importantly, they give us access to the tools and the methods and the people and their knowledge that are in these centers. Let me give you a few examples of how we have used them.

I'll first start with the aluminum beverage can, which you may think of as a kind of a convenient way to hold beer or pop, but it really is a very sophisticated pressure vessel. It has to be able to sit in somebody's trunk and be heated up to around 100 degrees and hold the beverage without popping. You need to be able to stack 25 cases of beer on top of each other and not have the bottom one collapse, and it has to have enough damage tolerance so that it isn't dented in handling.

Now there are something on the order of 10 to 15 parameters in a can design. That will give you something on the order of a million design possibilities, and we used to sort these by trial and error. Today we design beer cans and pop cans on a Cray in 1/20th of the time that it used to take us to do that. Our customers can walk in the door of the tech center and walk out with a can designed to their needs and specifications.

The second example is the aluminum-intensive vehicle which is a whole new concept of how to build the structural part of a car, and it's a switch from, of course, steel to aluminum, and so in order to design the car and assess the crash-worthiness of the components and assemblies of this car, we started developing math models, and they have allowed us to really design now a car that's probably the safest and most fuel-efficient lightweight, environmentally-sound car in its size in 1/10th to 1/20th of the time that it would have taken us to do it by previous technique, if we could have done it at all. So that was, again, using the supercomputing center in Pittsburgh.

And then, finally, that car concept is made up of extrusions and castings, and so we've developed a whole technology of how to design casting dies, how to design extrusion dies, so they run right the first time, and we've also been able to optimize the extrusion conditions and the casting conditions, and all of this is done in software and it's all been done on supercomputers.

In material science, microstructure is everything. It determines the performance of a piece of metal in any kind of application, and a microstructure is determined by the processing path that produces it: casting, forging, rolling, extrusion, these kinds of things. We have modeled materials from the atomic level to the size of cars, and so we now have a firm grasp on this continuum of processing microstructure product performance so that we can predict properties and optimize processing paths.

The second point I want to talk about is the HPCC Program, and I think it really is on track to deliver the critical computation capabilities that we need. Early access and exposure to these kinds of capabilities will allow industries like mine to increase our competitiveness. It allows us to increase the scale and the complexity of problems, and it also allows us to solve those problems in a very short length of time, which is a critical component of competitiveness today.

I think, more importantly, it allows our engineers to redefine the problems. It broadens their horizons so they envision a problem now much differently and from a totally different perspective than they did without access to this kind of computational capability.

The Grand Challenge programs in the universities and the government labs I think allows industry to hire scientists that have training and experience in high performance computing because we don't have time to train them ourselves and it is very important that we have access to a supply of engineers that have had experience using high performance computing.

Also, I think that there is sufficient evidence—and others could comment perhaps more than I could—that the HPCC has helped the supercomputing industry in the United States maintain its leadership.

And I guess the final point I want to make is to comment on the questions raised about the program focus, in particular, the issues raised about massive parallel processors. Now our problems are really very large-scale deformation, structural, mechanics problems and fluid dynamics problems that are associated with all aspects of the process of digging up dirt and turning it into aluminum cans and cars. The complexity of the models that we use to optimize and to design new processes is limited by computing power, and the lack of software for high performance architectures is really a critical obstacle in our ability to access and solve these problems. And this is really precisely where I think we ought to be spending our money, is to create the software.

We, Alcoa, has launched into a program using CRADA approach with Lawrence Livermore and Sandia to develop a finite element program for metal-forming analysis that is based on parallel processing. And our aim is really simply to solve our production problems in a much shorter time frame than we do today.

We are making a substantial commitment in time and in dollars and in technology because we believe the MPP architecture is the best answer. Now not all companies can necessarily have the resources to do this, and so, therefore, I think that the HPCC is playing a very major role in allowing this to happen. And, very specifically, I believe that HPCC should have an explicit objective to coordinate and further the development of scalable application software that takes advantage of the range of HPC hardware architectures and then this makes HPC available to a much broader segment of U.S. industry. I think this should include programming models and standards for the high performance computing architectures. It should promote support and encourage the development of parallel software environments, so they become more widely used and that they should be transportable from distributed clusters to massively parallel machines.

I think all of this will eliminate the barriers that exist out there today to making use of MPP machines and really allow a whole lot of industries to realize the full potential of this very powerful tool. Thank you.

[The prepared statement of Dr. Bridenbaugh follows:]

High Performance Computing and Communication
Testimony to the House Subcommittee on Science
1993 October 26

Mr. Chairman and Members of the Subcommittee:

On behalf of the Aluminum Company of America, I appreciate the opportunity to speak to you today about the value of the government's High Performance Computing initiative to Industry. My name is Peter R. Bridenbaugh, and I am Executive Vice President of Science, Engineering, Environment, Safety and Health for Alcoa. I was also a member of the NSF Blue Ribbon Panel on High Performance Computing, whose report was issued earlier this month. I would like to focus my comments today on how Alcoa is using HPC technology, and how we have benefited from having this general capability available to us, and then I'd like to comment briefly on the HPCC program itself.

The role of the NSF Supercomputing Centers in HPC technology transfer

Because most of Alcoa's experience with HPC has been enabled by the existence of the NSF Supercomputing Centers, I would like to first speak about the important technology transfer role that the centers have played. Alcoa has been an industrial affiliate of the Pittsburgh Supercomputing Center (PSC) since 1987. At the Alcoa Technical Center, located about 25 miles east of Pittsburgh, we perform research and development for all phases of aluminum production. Today, advanced research in materials requires the use of high performance computers to understand how materials behave under various conditions, how products can be better designed and how processes can be improved. As an example, we have used the PSC computers to design a "better" beverage can - one that uses less metal and thus is less costly and more energy efficient to produce. We are also using the computational power at the PSC to design new products. For example, we have used the PSC computers to estimate the crash worthiness of aluminum automobile components, which will ultimately lead to safer, lighter, more fuel efficient cars. Aluminum automotive products are a key part of Alcoa's future, and scientific modeling on the supercomputer helps us understand how to design these parts and how to design the manufacturing processes that will make those parts. We have used the high performance computing facilities at the PSC to analyze the detailed behavior of aluminum at the atomic level through final product performance, in order to improve our basic understanding of our materials.

Access to high performance computing facilities like those at the PSC helps us in two ways: it allows us to solve problems which are just too large for our in-house computers, and it allows us to solve time critical problems in a competitive manner. The latter ability is particularly important in the design of automotive components, where proposals often must be prepared within tight time constraints. Our access to a large computing facility helps to keep us competitive with our international competition. A computing resource like the CRAY is simply too expensive for Alcoa to purchase outright. Neither could Alcoa make full use of computers of this size. A communally shared resource, like those available at the NSF centers, is the ideal means of providing access to state of the art computing to American industry. Since they are committed to keeping their facility at the leading edge in the use of massively parallel computing systems, we are able to take advantage of the highest performing computing available, which allows us to explore ever more complicated material research issues.

The NSF Supercomputing Centers provide more than just access to machine cycles. As one of the explicit roles spelled out for the NSF in the HPCC program, the NSF through the centers serve as toolmakers for scientific programmers everywhere. They have the experience to recognize general needs of the scientific community and sufficient expertise to satisfy them. For instance, NCSA Telnet (developed at the NSF center at the University of Illinois) runs on virtually every Macintosh computer at Alcoa, and is the basic communications link from desktops tc many shared resources. Our affiliation with the PSC has benefited Alcoa's research efforts in other ways. We have used the PSC as a testbed for computing technology, and interaction with the researchers at the PSC has allowed us to benefit from their experience when installing new hardware and software.

In summary, the NSF Supercomputing Centers provide an efficient and productive way for industry to gain access to HPC technology. They provide access not only to actual machine cycles, but also to tools, methods, software, and, perhaps most importantly, people as technology transfer agents. This is a critical role that the government has enabled, and one that should not be overlooked.

The HPCC Program and the role of the Government

I would now like to turn my attention to the HPCC program itself. On a somewhat philosophical level, I would like to speak to why I believe that the government should have an active role in sponsoring and coordinating the development and dissemination of HPC technology. In my view, the critical long-term value of having a well-integrated, government-coordinated program in HPC is threefold:

- Increasing the productivity and competitiveness of American industry in general. This is only possible if industrial researchers and engineers have access to HPC facilities and have broad and early exposure to HPC capabilities. Industry will need scientists and engineers who can formulate problems to take advantage of the newer, faster architectures, such as MPP, and it takes time to understand these capabilities well enough to conceive of and structure problems appropriately. This is precisely why Alcoa got involved early with PSC. Our early use of the CRAY was mostly spent by researchers learning to think more broadly about their problems, for example, to begin to envision forging dies or extrusion problems in 3 dimensions. Broadening your toolset and understanding enough about the new tools to use them effectively takes time, and we must be sure that we have the time to become masters in the use of these tools, in order to maintain our competitiveness.

- Maintaining the US leadership in the HPC industry. This is an important industry segment, and one in which we have been able to maintain our leadership, due in no small part to the government's action in passing the HPCC Act in 1991.

- Grand Challenge research in universities and Government Laboratories. The Grand Challenge research benefits industry in many ways. It provides significant advances in fundamental scientific and engineering research, and is also fundamentally changing how research is performed. HPC availability in academia means that Alcoa is better able to hire engineers and scientists who have training and experience in the use of high performance computing in materials research. The Grand Challenge research and the resulting collaborations also result in the development and dissemination of advanced software methods and algorithms which can take advantage of parallel architectures.

Each of these long-term benefits I have just described have been furthered by the High Performance Computing and Communications Act of 1991. Although it is certainly true that the job is not finished, we have made considerable progress and we are beginning to see payback from our investments. From my perspective, we are clearly on the right track with this program, and we should maintain our pace.

Although I am not in a position to comment directly on some of the specific questions regarding the planning process or the role of the National Coordinating Office with respect to the HPCC, I would like to address some of the issues that have been raised about program focus, in particular, the questions that have been raised about Massively Parallel Processors (MPP).

I have heard the criticisms that the commercial applications of MPP are nil, that MPP systems have not proven effective, and that MPP systems are too hard to program. The number of successful commercial applications of MPP today may be small, and there are indeed barriers to using MPP systems effectively, but these are precisely the areas where I believe we should be focusing more effort. Within Alcoa, our computational problems are generally large deformation structural mechanics and fluid dynamics problems associated with chemical and aluminum processing. We have seen a tendency for problem sizes to grow at about the rate of growth of processor speed. In our current environment, researchers must reduce the complexity of their models to accommodate the size of the machines available. Increasing compute power allows us to solve current problems with better accuracy (because we can use a more complex model), and to attempt problems that we currently do not. Examples of this latter kind include all but the most rudimentary extrusion models, most 3D forging models, and large die casting problems. Access to more powerful processors can fundamentally change the way we approach our problems.

Because Alcoa recognizes the potential benefits of MPP and because we also recognize the lack of software for high performance architectures as a critical obstacle, we are developing a parallel finite element program for metal forming analysis in collaboration with Sandia and Lawrence Livermore Laboratories (CRADA 92-0324 "Precision Aluminum Forming") It is our intent to use the results of this CRADA to give us access to the current generation of massively parallel processors for our difficult production and research problems. If we can solve large production problems we can reduce time to market, and improve product quality. For instance, currently it is a common occurrence for the first design of an extrusion or forging die to fail to produce parts which are in specification without time consuming experimentation and modification. With the capability to perform sufficiently detailed analyses, we can increase the percentage of first time die successes, resulting in significant savings in both time and money.

We are willing to make this considerable investment in time and dollars to develop this advanced software capability in order to take advantage of the MPP architecture. Many industries are not able to make this investment. Ultimately, Alcoa, and the rest of industry, would prefer that this type of application software be commercially available, off-the-shelf, and able to run on a wide variety of architectures, from workstation clusters to supercomputers. Commercial software is unlikely to become widely available on high performance architectures unless a more coordinated effort is undertaken.

Therefore, I believe that the HPCC program should have an explicit goal of furthering the development of commercial application software for scalable HPC machines. This goal should be supported by the development of programming models and standards for HPC architectures. The HPCC program should work to promote, support and encourage the development of parallel software environments so that they become standard and widely available. This will allow the development of commercial codes and will also enable end users to write their own software. It's essential that software be portable from distributed clusters to massively parallel machines. This change in the direction of the HPCC will eliminate the barriers to effective use of MPP systems, and will allow more of industry to realize the potential of this significant processing capability.

75-413 O - 94 - 4

Summary

Let me close by briefly summarizing my key points.

- The NSF Supercomputing Centers have served an important role in transferring HPC technology from the universities and government laboratories into industry, and continue to do so through their charter in the HPCC program. This is a critical function that should continue.

- The HPCC program that was established in 1991 is on track to develop and deliver to American industry critical computational capabilities that will be required to remain competitive into the next century. Early access and exposure to HPC capabilities will enable industry to increase the scale and complexity of their research, and also to fundamentally change the way that research is performed.

- The HPCC program should have an explicit objective to coordinate the development of commercially available scalable application programs that take advantage of the range of HPC hardware architectures, so that the benefits of HPC can be shared by a broad segment of U.S. Industry.

Thank you.

Dr. Peter R. Bridenbaugh
Executive Vice President of Science, Engineering, Environment, Safety and Health
Aluminum Company of America
Alcoa Technical Center
100 Technical Drive
Alcoa Center, PA 15069-0001

Examples of HPCC Usage in Alcoa

Beer Can Design

Beverage containers have three major performance requirements: they must be able to hold a certain internal pressure without deforming (so full cans don't deform when the contents are heated), they must be able to withstand a drop of a certain height without deforming (so normal handling doesn't cause damage), and they must be able to support a certain amount of column load (so the operations of putting on the top of the can and filling it don't cause damage). We are able to model the performance of a proposed can design quite accurately and evaluate the quality of a design without manufacturing costs and delays. There are generally about 10 to 15 design parameters (e.g., can diameter, dome radius, and base diameter,) that impact the performance of a can design. The process of design is often one of trial and error, in a design space of millions of plausible can designs. We have used the PSC's CRAY to generate and evaluate hundreds of designs automatically, and to automatically cull the good designs for further analysis. This work took one twentieth of the time it would have taken on our in house computers.

Bumper Design

We have used the PSC's CRAY to analyze bumper designs in one tenth to one twentieth the time we could perform the analysis on our in-house computers. This allows us not only to respond to the tight time constraints associated with bidding to our automotive customers, it allows us to evaluate more alternatives in a short time (which often results in better designs). On several occasions, had we not had access to the PSC CRAY we would not have been able to respond with a timely bid, which would have jeopardized our ability to compete in this market. For example, we might get a request from Detroit that they need a bumper for a light utility vehicle that would dissipate energy to a certain specification in a 30 mph crash, and that they need the design by next Thursday to meet design schedule. The availability of the CRAY makes the difference in whether we can compete or not.

Automotive Die Casting Analysis

When casting complex parts it is useful to know how the molten aluminum flows to fill a part. It is important to avoid the formation of vortices (which are likely to trap gas), to avoid the formation of weld lines in the finished part (which can cause subsequent weakness in the part), and to ensure that the first material to enter the die sweeps the impurities and oxides out of the part. When dies are designed without analysis, part defects can cause expensive and time consuming die corrections. In addition, it can take months of trial and error experimentation to determine the correct casting practice to manufacture good parts.

At the Alcoa Technical Center we are using modeling to decrease the number of die corrections required, and improve the turnaround to produce new parts. Currently we can model the mold filling of small (smaller than a breadbox) parts in a matter of days to weeks of computation on a fast workstation (a few days of computation on a CRAY C-90). This ability has been used to good effect in our automotive program, where we have had some first time die successes as a result of computational modeling. Unfortunately, many of the parts we want to model are larger than we can solve on our current facilities. When this happens we must substructure the large parts into pieces we can handle, and make some educated guesses as to the appropriate way to fit the parts together. To solve the larger automotive part mold filling problems will require at least an order of magnitude increase in computational capability (both size and speed).

After the mold has been filled the parts must solidify. Since metal shrinks as it freezes, it is important to ensure that unconnected liquid pockets are not formed in the process of solidification, or porosity will result. While we can solve part of this problem approximately, the underlying thermo-mechanical analysis will require at least a two order of magnitude increase in computational power before we can attack it for production parts. These analyses would have important ramifications not only for shrinkage porosity, but also for subsequent heat treating.

Automotive Extrusion Modeling

Extrusions are critical components in several automotive programs. When new extrusion components are designed, it is often necessary to experiment with extrusion press operations to determine how to cool the extrusions to avoid distortion of the component on quenching. We have devised a numeric model of the extrusion quench process. Computer runs of approximately 30 hours on a fast workstation are required to simulate 6 inches of extrusion. To determine extrusion quench process parameters at least several feet of extrusion must be simulated. For this analysis to be performed on a production basis would require a computing facility at least 2 orders of magnitude faster/larger than is currently available.

Extrusion dies often do not perform correctly on the first design iteration. Our preliminary modeling efforts using two dimensional techniques look promising, but both our computing facilities and algorithms are currently inadequate to solve the extrusion models for die design in three dimensions.

Fundamental Understanding of Material Behavior

Our customers today, including automotive, aerospace and beverage can manufacturers, demand materials with enhanced performance requirements like strength, stiffness, formability, and with new performance hurdles, such as maintainability, recyclability, and predictability. Meeting these demands requires a profound understanding of not only our own materials, but of our customer's specific design requirements. This understanding can only come through detailed models that quantitatively describe the relationships between how we process aluminum and the resulting microstructure, that is, what the material looks like at the atomic level, and models that describe how that microstructure affects final product properties. This gives us the information we need to optimize and control final product performance.

As an example of this type of fundamental material modeling, we have studied the behavior of aluminum crystals. Aluminum is a crystalline material. On a single crystal level, the mechanical behavior of aluminum is radically different than the behavior of a chunk of aluminum containing many thousands of crystals. The single crystal exhibits pronounced anisotropy - how stiff it is depends greatly on the loading direction. This radical anisotropy tends to get averaged out in a macroscopic chunk of material, but not completely. The macroscopic anisotropy of aluminum is due to the non-uniform distribution of crystal orientations in the sample, which is referred to as crystallographic texture. This distribution of orientations is affected by the processing of the material, in ways that we are still trying to understand. It is essential to quantify and control the anisotropy of any aluminum product that is going to undergo subsequent forming (like beverage can sheet).

At the Alcoa Technical Center, .'e have modeled two dimensional collections of tens to hundreds of aluminum crystals, in an attempt to describe this evolution of texture in a quantitative manner. These numerical experiments ran in a couple of days of CPU time on a CRAY Y-MP, and were compared .o results from physical experiments. The study indicates that fully three dimensional models are necessary to simulate the evolution of the crystal orientation distribution adequately. In order to carry these computational experiments to three dimensions and thousands of crystals, increases in computer capacity of at least two orders of magnitude will be required.

Software Development at the NSF Centers

ImageTools, developed at the San Diego Supercomputer Center) is used by researchers at ATC to convert among various file image formats. GPLOT (developed at the Pittsburgh Supercomputer Center) is the common CGM display software at ATC, and virtually everywhere else on the Internet. Simply stated, GPLOT provides a way to transfer graphic images. Personnel at the NSF centers also modify existing software from other universities to suit their computing environments, and these modifications quickly are incorporated into the basic releases. For instance, the PSC has modified and enhanced the Distributed Queuing System (DQS) software from Florida State University. DQS is a software system for exploiting interjob parallelism. That is, DQS is used to automatically distribute a set of jobs across a collection of machines in a rational manner. From the perspective of the users submitting the jobs, the collection of machines can be viewed as a single resource for submitting jobs. AT ATC, we are beginning to use DQS to take advantage of unused computational cycles on our workstations and desktop machines.

Mr. BOUCHER. Thank you very much, Dr. Bridenbaugh.

Dr. Rubbert.

Dr. RUBBERT. Mr. Chairman, members of the subcommittee, I'm the Unit Chief of Aerodynamic Research for Boeing Commercial Airplane Group. I'm one of the people in our industry who has led the revolution from wind tunnel-based aerodynamic design processes to processes that rely very extensively on advanced high performance computational capabilities. Today those computational design processes stand alongside the wind tunnel in terms of their importance to what we do.

The principal responsibility of my group is to produce the computational aerodynamic capabilities that Boeing uses to design its commercial jet transport airplanes. Those computational capabilities are extensive and are run in excess of 15,000 times a year by our design engineers. We utilize a range of computers from desktop models to the largest Cray supercomputers, and we work closely with our Boeing computational laboratory which has a massively parallel machine and a variety of other experimental computing engines.

I've also been involved in NASA's computational aerosciences project within the HPCC program, as a member of a NASA Advisory Council Task Force on Advanced Computing Capabilities, and I am a member of a joint industry-academia-NASA working group that undertook to replan the HPCC computational aerosciences project.

I would like to focus my remarks today in the area involving industrial competitiveness. One of the explicit goals of HPCC is to "spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process." Those processes reside in and are controlled by industry. Therefore, it is imperative that the industry be participants in the upfront program planning as well as the execution. They are the real customers.

The fact that the nongovernment HPCC Advisory Committee was never implemented, even though it was authorized in the 1991 Act, was, therefore, a mistake. In NASA's computational aerosciences project, we recovered from that mistake by implementing a program replanning process that involves strong and direct participation of industry and academia working hand-in-hand with NASA. That process is working very well.

We successfully arrived at a set of goals and an overall program description that was endorsed by industry, academia, and NASA as supporting the federal goals of providing broader benefits to industry and academia and effectively exploiting our differing core competencies. That is the type of result we want to achieve, and I strongly commend NASA for the speed, willingness, and sincerity by which they entered into that new planning process. They really did it right—and it entailed a culture change for them and for industry and academia, as their role shifted in the direction of becoming a supplier and facilitator.

NASA's computational aerosciences project today is clearly aimed at accelerating the development and facilitating the adoption of high performance computing by the U.S. aerospace industry. That is good. The question is, will it happen? HPCC is strongly focused on massively parallel architectures and tera flop broad performance. The aerospace industry needs are for computers that offer a balanced blend of affordability and superior price performance; ease and speed of programming; short total cycle times for problem execution; mature, reliable systems and systems software; adherence to widely accepted standards; that are available as distributed systems rather than large, data center sited installations, and that have some ability to solve very large application problems. We need computers that offer a balanced blend of those attributes, and if massively parallel machines turn out to do that better than any other, then we will acquire and use them to increase our productivity and competitiveness. If they don't, we won't. That represents risk.

In my opinion, HPCC is probably too narrowly focused on massively parallel computers and in trying to force the structure of application problems to fit the architecture of those computers. We need also to push the frontiers of computers that may not be massively parallel and which may not promise to achieve terra flop performance at any cost, but which do an overall better job of meeting the needs of industries which use high performance computing to improve their design and production processes.

Mr. Chairman, that concludes my prepared remarks.

[The prepared statement of Dr. Rubbert follows:]

Statement of
Paul E. Rubbert
Unit Chief, Aerodynamics Research
Boeing Commercial Airplane Group

Hearing on the High Performance Computing and Communications Program

Subcommittee on Science
Committee on Science, Space and Technology
United States House of Representatives

October 26, 1993

P. 1
Statement of P. E. Rubbert
High Performance Computing and Communication Program

Mr. Chairman and Members of the Subcommittee:

My name is Dr. Paul E. Rubbert. I am the Unit Chief of Aerodynamics Research for The Boeing Commercial Airplane Group. A principal responsibility of my group is to produce the computational aerodynamics capabilities and codes that Boeing uses to design its commercial jet transport aircraft. For example, the aerodynamic design of every portion of the exterior surface of the new Boeing 777 transport aircraft involved the use of these computational aerodynamic codes, with the role of the wind tunnel being principally one of design validation. We employ a variety of computers in our work, ranging from desk top models to a large Boeing-owned CRAY Y-MP supercomputer. We access the large NASA computers that comprise the Numerical Aerodynamic Simulator (NAS), and we work closely with an adjacent Boeing computational laboratory which has a modern massively parallel computer and a variety of other experimental computing engines.

I have been involved in the NASA portion of the HPCC program as a member of the NASA Advisory Council Task Force on Advanced Computing Capabilities that produced a report on NASA's High Performance Computing and Communications Program in June of 1993. Subsequently, I worked closely with NASA as a member of a combined industry/academia/NASA working group on revising the plan for the Computational Aerosciences (CAS) project of the NASA HPCC program.

My principal focus has been on the CAS project of the NASA HPCC program, and the majority of my remarks today concern that project.

REVIEW AND EVALUATION OF THE CAS PLANNING PROCESS AND RESULTING PLAN

The Computational Aerosciences (CAS) project of the NASA HPCC program has served as a major catalyst for positive change in NASA's planning process. The agency, on information from the Advisory Council Report of the Task Force on Advanced Computing Capabilities (June 1993), modified its initial program planning process and plan to more effectively capture inputs from industry and academia, and to better exploit the differing core competencies of industry, academia and NASA. This is leading to the development of well-defined and complementary roles for each sector, and to the development of a one-team spirit among industry, academia and NASA.

NASA achieved this by calling together representatives of the major U.S. airframe and engine companies and academic leaders to solicit their advice and support in replanning the CAS project of the HPCC Program. This step led to the creation of a smaller industry/academic/NASA working group, and enabled industry and academia to have considerable input to the replanning process.

This working group subsequently produced an executive summary description of a revised CAS project that was endorsed by industry, academia and NASA participants as supporting the goals of the HPCC program, providing broader benefits to industry and academia, and effectively exploiting the capabilities of each sector.

To date, however, a detailed program plan capturing the intent of the executive summary has not yet emerged. While initial drafts of the plan were judged by some industry participants as not fully reflecting the summary intent, NASA is currently working to address this concern.

I am pleased and gratified with the speed, sincerity and dedication that NASA has shown in moving toward greater industry and academic involvement and participation. The agency's actions display a focus on supporting industry's needs, and an awareness of the complementary roles of industry, academia and government.

BENEFITS TO THE U.S. AEROSPACE INDUSTRY

The CAS Project has as its specific goals:

- To accelerate the development and availability of high performance computing technology (hardware, systems software, and applications software) that will be of use to the United States aerospace community.
- To facilitate the adoption and use of this technology by the U.S. aerospace industry.
- To hasten the emergence of a viable commercial market for hardware/software vendors to exploit this lead.

These goals, if met, will accomplish the overarching federal HPCC goal of spurring gains in the productivity and industrial competitiveness of the U.S. aerospace industry.

The HPCC program and the CAS project are clearly focused on accelerating the development and availability of massively parallel computers and computing, and in my judgement will be effective in accomplishing this. It will hasten the time at which we develop a better and clearer understanding of the role or roles that this type of computer can have in enhancing our engineering design processes, and it will hasten the availability of the systems software necessary to support code development and applications.

The unanswered question at this point is to what degree this will be of value or use to the United States aerospace industry. Our needs are for cost-effective computational capabilities that support improvements in the processes we use to design and produce airplanes; with improvement defined in terms of reduced process cycle time, reduced cost and reduced variation. The paradigm shift from wind tunnel to computation-based design processes is now mostly behind us, and we have an increasingly clear vision of the computer attributes that best support our present and future needs. Those attributes include the following:

- Superior price/performance, to support our need to reduce costs.
- Speed or ease of programming, to support just-in-time application code development.

- Small overall flow time to execute problems, to support our requirement for reduced cycle time.

- Mature systems that work as advertised, with high reliability and good system software.

- Adherence to widely accepted standards, to ease the problem of software portability.

- Available as distributed systems that can be under the local control of engineering process owners. Our extensive experience with centrally managed, data-center-sited supercomputers is leading us away from that type of configuration as a preferred host.

- Some ability to solve very large problems. Our application needs span the range of problem sizes from small to the very largest. However, since the flow time for development, engineering validation, and engineering process acceptance of new application codes is measured in years, we find that by the time an application code has passed all of the tests and gains acceptance for use within an airplane design process, that code is generally executable on less than the most powerful available computing host. This means that we will not, in general, need the latest-and-greatest supercomputer to support the day-to-day work of airplane design and the continued improvement of our design processes. We do need some access to the largest computers to carry out research or to occasionally host an unforeseen, very large problem that may arise, but that need will not dominate our future computer acquisition decisions, particularly if a National facility such as the NAS is available to support computing-consumptive research or to host the occasional unforeseen very large application.

Now let us return to the unanswered question concerning the value of HPCC to the United States aerospace community. The answer is that if the massively parallel computers that arise from HPCC activities turn out to be the most effective among all candidate computers, parallel or not, as measured by a balanced assessment across the various attributes that are important to us, then they clearly will be of value and we will acquire them. HPCC will have facilitated their adoption and use by industry, will have hastened the emergence of a viable commercial market for the hardware/software vendors, and will therefore have met the stated goals of the CAS project.

Whether that will occur remains to be seen. At present, massively parallel computers appear advantageous only in certain "niches" where the nature of the problem is particularly amenable to massively parallel processing. The major challenge is to find ways of solving the application problems of interest that effectively exploit a massively parallel architecture, and the CAS project addresses that challenge.

This naturally raises the question of are we directing too many of the HPCC resources into the development of massively parallel computers? I think that the answer is probably yes.

What I see happoning is a strong attempt to transform the structure of application problems into a form that can be dealt with efficiently with a massively parallel computer architecture. Do we have the tail wagging the dog? An alternate approach would be to place more emphasis on developing computer architectures that are best suited to solving the application problems hat need to be solved, rather than the other way around. The present vision of my technical staff is that modestly parallel computer architectures, involving tens rather than thousands of processors, are better suited to solving the types of multidisciplinary design and optimization problems that characterize airplane design and our future needs.

I am also concerned about the risk being imposed upon our computer vendors because of the very strong focus on massively parallel. What if massively parallel machines fail to measure up to industry's requirements, and industry ends up going down a different path and purchasing other types of computers? Those vendors who chose to gamble everything on the massively parallel rainbow will find that the market is not large enough to provide an adequate return on investment, and they will have forfeited their leadership position.

Another aspect related to benefits for the United States is the competitive situation with other countries. In the commercial transport aircraft industry our foreign competition is Airbus. Since our computer vendors sell their products worldwide, Airbus will be able to acquire and exploit the same computers. Our competitive advantage will come from doing the right things sooner and better, namely making earlier and more informed decisions concerning which computers to acquire and when, and being able to develop and exploit applications software more rapidly. Industry's active involvement in the CAS project will posture us to capture those competitive advantages.

DOES HPCC REPRESENT AN APPROPRIATE FEDERAL ACTIVITY?

The breadth and impact of the federal HPCC program spans the worlds of both science and engineering, which is a very broad spectrum. The support of science objectives, education, and of national security have been appropriate and traditional roles of the federal government. Those roles are clearly imbedded in the HPCC program in the form of fostering development of the world's most powerful computers, which are the tools of choice for (i) the scientist who seeks solutions of problems that could not be solved before, and (ii) the guardians of national security for whom the cost of a supercomputer is inconsequential compared with the alternative of possibly ending up second best in an international confrontation.

A federal effort to "spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process" is less traditional for the United States. However, we find ourselves in a world economy wherein our trading partners employ selective government involvement in support of their private industries. Therefore I think that the HPCC program should seek to take advantage of the opportunity to spur U.S. productivity and industrial competitiveness.

However, we are learning our way into how best to do this, and we have to be cognizant of the differences between doing "what's right for science" and "what's right for engineering in a competitive commercial environment." In the pursuit of science objectives, a successful approach has been for a research-oriented organization such as NASA to play a singular lead role in setting the objectives and laying out a program plan. It assumes the role of the "customer", and other organizations which assist in executing the plan assume the role of "supplier". But that is not the right role for a research-oriented organization when supporting engineering objectives in a competitive commercial environment. In that environment, industry is the "customer" and NASA is the "supplier", which is nothing less than a role reversal!

In the case of NASA's CAS Project replanning activity, this led to the U.S. aerospace industry adopting the role of the customer. During the replanning activity, it was advocated that NASA would be more effective by playing the following role:

1. NASA would play a strong role in facilitating a strong, direct, and continuous broad-band communication channel between the computer industry, the aeronautics industry, and academia. These linkages, and the types of information that must flow between the computing and aeronautics industries and academia are shown in Figure 1. The impact of this will be that:

   - Our computer vendors will understand the emerging computing requirements of our aeronautics industry, thus enabling them to develop advanced computer architectures that meet those requirements.

   - Our aeronautics industry will have a clear and continuous view of advanced computing products that are upstream in the pipeline. This will enable them to make early and informed decisions on how to best utilize new advances in computing in ways that will enable them to capture gains in productivity and competitive advantage.

   - Academia will once again be able to function effectively as a supplier of advanced algorithm technology for the new and emerging computer architectures.

2. Funding for academia will once again be available from NASA. Up until the early 1980's, academia was a major supplier of key elements of computational algorithm technology. Most academic researchers were supported by NASA.

3. NASA contract funding for algorithm research and code development will once again be available to researchers within the aeronautics industry, thereby accomplishing the following:

   - Eliminating the "technology transfer" time gap. Industry exploits advances in algorithm technology or other forms of computing technology years earlier if their own researchers participate in its development.

P. 6
Statement of P.E. Rubbert
High Performance Computing and Communication Program

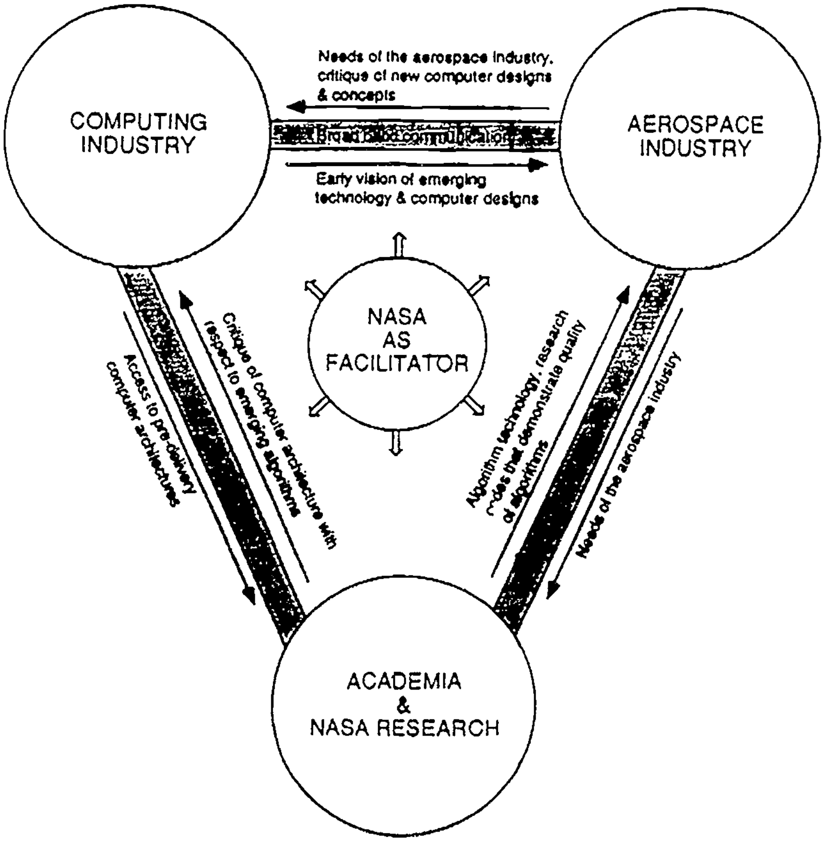

Figure 1: Key Ingrediants of a Successful CAS HPCC Program

- Assuring a greater degree of industry relevance.
- Enhancing industry's skills, knowledge and training in algorithm research. Those skills are extremely important in enabling industry to effectively engage in the continuous improvement, extension and enhancement of its codes.

A major activity of the CAS replanning effort was to move NASA away from its customary role as a strong singular leader and solitary participant and toward the role outlined above. I feel that we achieved a very significant accomplishment.

In summary, I believe that HPCC indeed represents an appropriate federal activity. It is very broad in span, and includes both traditional and less traditional roles of the federal government. Across this broad span it is important that the differences between science and between engineering in a competitive economic environment be recognized and accommodated. That accommodation should be reflected by the adoption of modified roles of the involved federal agencies.

LACK OF A NON-GOVERNMENTAL HPCC ADVISORY COMMITTEE

The 1991 act authorizing the HPCC program created a non-governmental advisory committee to provide the OSTP director with an independent assessment of the HPCC program, to provide the HPCC program with access to expertise from outside of the federal government and to ensure that the planning and implementation of the HPCC program incorporated the views of groups which will be R & T performers, suppliers of services, and users of networks.

To date, the HPCC advisory committee has not been put in place. This raises the question of whether the non-federal segment of the R & T community has been adequately represented in the development of the HPCC program goals and implementation plan.

The answer is that in a broad, far-reaching program such as HPCC which endeavors to "spur gains in U.S. productivity and industrial competitiveness by making high performance computing and networking technologies an integral part of the design and production process," it is imperative that the industry be involved in the development of program goals and implementation plans. This was made clearly evident in the CAS replanning activity, which demonstrated the essential role that must be played by industry. I feel that the CAS replanning process could serve as a role model for other elements of HPCC that may not have benefitted from a proper amount of industry involvement.

The process of developing implementation plans must contain a proper recognition of and respect for the differing core competencies of government research laboratories, industry and academia, and must seek to exploit the best of each in complementary roles. A planning approach that involves representation from each of these segments also leads to a one-team-spirit and harmony between the various participants. It is a fact of life that in times of shrinking budgets, absence of any of these participants in the planning process will result also in their absence from participation in the execution of the plan.

P. 8
Statement of P. E. Rubbert
High Performance Computing and Communication Program

IN CONCLUSION

The goals of the HPCC program are sound and represent an appropriate federal activity. The program has the potential of spurring gains in the productivity and competitiveness of the U.S. aerospace industry by favorably impacting our engineering design processes. But, that will only occur if massively parallel computer architectures turn out to be most effective among all candidate computers, parallel or not, as measured by a balanced assessment across the various attributes that are important to the aerospace industry. That remains to be seen. As a minimum, the HPCC program will provide us with an earlier understanding of the appropriateness, or lack thereof, of massively parallel computers in our future. This leads me to believe that we are perhaps putting too many of the resources into the development of massively parallel computers at this point, and that more emphasis should be placed on other architectures such as modestly parallel.

The CAS replanning activity clearly demonstrated the absolute necessity of getting industry involved in the initial planning of activities such as HPCC that aim to influence or impact the design or production processes in use by industry. That type of industry involvement and participation can require a culture change and role reversal for federal research laboratories that are accustomed to the role of singular leader and goal-setter in a science-oriented environment. The NASA is to be strongly commended for the speed, willingness and sincerity by which they learned how to effectively capture and utilize the insight and perspective that industry and academia contribute to the program planning process.

P. 9
Statement of P. E. Rubbert
High Performance Computing and Communication Program

Paul E. Rubbert
Unit Chief, Aerodynamics Research
Boeing Commercial Airplane Group

Dr. Rubbert leads the group at Boeing that produces the computational aerodynamics capabilities and codes that Boeing uses to support the process of designing its commercial jet transport aircraft. Those capabilities are extensive, and widely used throughout Boeing.

He is a pioneer in the field of Computational Fluid Dynamics who led the revolution from wind-tunnel-based airplane design processes to processes that rely extensively on advanced computational capabilities. Codes produced by his group are executed more than fifteen thousand times annually by Boeing design engineers, and they influence the aerodynamic design of every external portion of Boeing transport aircraft. His group employs a variety of scientific computers ranging from desk top models to the large Boeing-owned CRAY Y-MP supercomputer to the array of NASA computers that comprise the NAS. And they work closely with an adjacent Boeing computational laboratory which has a modern massively parallel computer and a variety of other experimental computing engines.

Dr. Rubbert is a graduate of M.I.T. in the field of aerodynamics. He has served on many governmental advisory committees and panels. His honors and awards are extensive and include his election to the National Academy of Engineering, election as a Fellow of the American Institute of Aeronautics and Astronautics, and election as a Technical Fellow of The Boeing Company.

Mr. BOUCHER. Thank you very much, Dr. Rubbert.

Dr. Frazer, we'll be pleased to hear from you.

Dr. FRAZER. Mr. Chairman and members of the committee, my name is Don Frazer. I'm Vice President of Massively Parallel Products for Oracle Corporation. Some of you may not be familiar with Oracle. We are the largest supplier of data management software and services in the industry. Over the past four quarters we've had over $1.6 billion in sales, over 60 percent of which came from abroad.

I've been involved with high performance computing of all varieties—scalar, vector, and massively parallel—with IBM and Teradata before I was with Oracle, for a number of years.

I appreciate the opportunity to appear before you today to present a perspective on massively parallel supercomputing from the point of view of someone who's been actively engaged in the development of commercial software and to refute some of the statements that have been made, particularly in written reports, about the relevance of massively parallel computing to our industry.

At Oracle we regard massively parallel supercomputing as not only highly relevant, but absolutely critical to our future success. We believe that this technology will be a key factor in this country's national efficiency and competitiveness.

If I could summarize just quickly five points, we think that the HPCC Program has made fundamental and critical contributions to our national capability and should be continued. We believe that massively parallel supercomputing is essential for a broad class of *nonnumeric* applications in addition to the numeric applications that have been most of the focus historically. We have devoted substantial time and energy and resource to massively parallel supercomputing and we're continuing to do so.

We believe that HPCC funding should shift from its historical focus on hardware to focus more on software and that HPCC funds should be allocated to purchase or enhance commercial off-the-shelf software, known as COTS, wherever possible.

Finally, we believe that HPCC funds should be expended only by competitive processes to ensure maximum return for taxpayer funds.

The last points are particularly important. If research and development is required for capabilities beyond those available in commercial off-the-shelf software, we believe that the money should go toward extending the capabilities of commercial off-the-shelf software by focusing on the leading- edge software industry in this country and funding only the incremental development necessary to achieve the required extensions to capabilities.

We strongly support Congressman Walker's amendments to H.R. 1757, which I understand were accepted by you, Mr. Chairman, and the rest of the committee unanimously, to require purchase of COTS wherever possible and to require competitive bidding on all awards over $25,000.

The rate of generation of computer data is skyrocketing in the commercial world. Scanners, mail, and telephone order merchandising, automatic teller and ticket machines, credit and debit card purchasing, and other sources are combining to produce enormous volumes of data which is presented to computers at prodigious rates which exceed the capabilities of the largest mainframe systems we have today.

It's been demonstrated repeatedly that the data that's generated contains a wealth of information which is of tremendous value to the enterprise that's capable of understanding and exploiting it, but, unfortunately, the largest mainframe systems have been incapable of generating more than expensive and inadequate summaries or statistical samples of the data. We believe that the solution to both the capture of this data, and the 'mining' of the information that resides in it, lies with massively parallel supercomputing technology. These supercomputers can add computing power, data storage capacity, and input/output capability with near linear gains over a very broad range which extends far beyond the capability of the largest mainframes today.

We have been adapting our relational database software to exploit massive parallelism for over five years, and we've spent to date over $11 million and we don't think we're done. We think we have a significant lead over our competitors, but we're continuing to invest because everything we've seen reinforces our conviction that we're on the right track.

You need only compare the list of vendors announcing or shipping massively parallel platforms today to the list of a year ago to see that we're not alone. The list has nearly doubled in the last year, and there are several other systems in the wings, vendors we're discussing porting our software to, which will enlarge it even more.

In addition, several of our software competitors have indicated by their announcements and actions that they share our vision. And finally, subsequent to the preparation of my written remarks I learned that there was a symposium conducted by MCC in Texas last week which was attended by more than 200 people. The subject of it was massively parallel computing's relevance to the commercial marketplace.

Relational database is but one area in which Oracle sees its future tied to this technology. We have a text search and retrieval product which we're porting to a massively parallel platform. We're engaged in a joint study with the Patent and Trademark Office. We see this as the absolute key to searching and understanding large bodies of text.

We have an office systems product which provides electronic mail, meeting scheduling, and a variety of other office functions. We have ported that to a massively parallel machine and we expect by the end of 1993 to have somewhere between 800 and 1,000 of our employees totally dependent on this machine for all of their office functions. Now this is in a company where you can't do anything without electronic mail. You can't order a paperclip or reserve a conference room or anything. So these people—we are really betting on this technology.

Finally, we've announced an agreement with U.S. West to perform a trial of electronic service to the home, including messaging, text retrieval, and video on demand, all of which will be based on massively parallel supercomputing technology.

I would like to turn just quickly to future funding targets for HPCC. I believe that the most important areas—and I share this view with many of my co-panelists, I think—are software, but one difference between the kind of software I'm talking about and the kind of software that characterizes most of the calculations that have been performed historically on massively parallel machines is that the latter have a beginning and an end. Typically, you run a computation; you stop the computation; you can reinitialize, service the machine, do whatever you want. But electronic mail and credit and debit processing and all of the things I've talked about above are not in that category. And so one of the major themes is that we need work on software to support five-to-seven-day-per-week, 18-to-24-hour-per-day operation on very, very large machines.

In my written testimony I have outlined five specific proposals for software research topics. The overall focus is in three subject areas. First is continuous operation in spite of failures and enabling service and upgrades of the hardware and software, and particularly the software, to occur online. The second is manageability of large numbers of processing systems working closely together, and the third is tools and techniques to make massively parallel software more easily accessible to the general population of software designers and developers, and then to the country at large.

Mr. Chairman and members of the committee, thank you for the opportunity to appear here to present a view of massively parallel supercomputing technology as seen through the eyes of someone in the commercial software development field. I hope that my comments will be useful. As beneficiary of past funding, we salute you and we look forward to a shift of funding from hardware to software.

Thank you.

[The prepared statement of Dr. Frazer follows:]

Written Testimony of W. Donald Frazer
before the
Committee on Science, Space, and Technology
Subcommittee on Science
U.S. House of Representatives
October 26, 1993

Mr. Chairman, and Members of the Subcommittee:

My name is W. Donald Frazer, and I am Vice-President, Massively Parallel Products, for Oracle Corporation, a California based company operating in over 100 nations throughout the world. Oracle is the largest supplier of data management software and services in the industry, with revenues in our last four quarters of $1.6 B, over 60% of which came from sales abroad. I have been involved with high performance computing–scalar, vector, and massively parallel–as a hardware and software developer with various companies for a number of years.

I appreciate the opportunity to appear before you here today to present a perspective on massively parallel supercomputing from the point of view of a successful developer of commercial software, and to refute some misconceptions put forth in recent assessments of the HPCC program's relevance to our industry. At Oracle, we regard massively parallel supercomputers to be not only highly relevant, but critical to our future success. Even though the commercial exploitation of this technology has just begun, we believe that massively parallel supercomputing is destined to have enormous impact on the way in which all large collections of data are recorded and exploited, and that it will be a key factor in maintaining and enhancing our national efficiency and competitiveness. In fact, we stand at the threshold of a period of broad acceptance and proliferation of this technology across industry and government for commercial purposes.

In brief summary, we believe that:

- The HPCC program has made fundamental and critical contributions to our national capabilities, and should be continued.

- Massively parallel supercomputing is essential for a broad class of *non-numeric* business applications, as well as for those numeric applications which have been the focus of most attention to date. We have devoted substantial and expensive corporate resources to exploit this technology, and are continuing to do so.

- HPCC funding should shift from its historical focus on hardware to concentrate on software.

- HPCC funds should be allocated to purchase or enhancement of commercial off-the-shelf software (COTS) wherever feasible, rather than to the development of redundant packages from scratch.

- HPCC funds should only be expended by competitive processes, to ensure that maximum value is received with taxpayer funds.

These last two points are particularly important for the HPCC program to meet its objectives responsibly. If research and development is required for capabilities beyond what is available in COTS, HPCC should begin by looking to cutting edge software companies, such as Oracle, to provide these extensions or enhancements, and fund only the incremental R&D costs required to meet those specific needs. Furthermore, such software should be "open"; it is wasteful to subsidize hardware manufacturers in the creation of proprietary systems.

We strongly support Congressman Walker's amendments to H.R. 1957, which were accepted unanimously by the Chairman and Members of the committee. They require purchase of COTS wherever possible, and require competition for any award over $25000. These are vital and necessary additions to the HPCC legislation.

Oracle's Commitment to Massively Parallel Technology

Oracle is a "server" software company. This means that our products are meant to serve the needs of groups of users, who access them through a variety of "client" desktop terminal devices ranging from limited function terminals at the low end, through PC's, to the highest performance workstations. Our server software products run on almost all important server hardware platforms worldwide, ranging in performance from large workstations and PC's through mid-range systems to mainframes and massively parallel supercomputers, all with a common set of application and end-user interfaces. It is important to us to be able to support a wide variety of choice in hardware platforms for our customers, so that those customers retain the greatest flexibility to choose hardware platforms to fit their needs, while maintaining maximum leverage from training and experience among their staffs. The customers also preserve an "open" environment, since we support the latest industry standards in our products.

Our products run today on, or are in the process of being adapted to, massively parallel supercomputers produced by nCUBE, IBM, Meiko, NCR, ICL Ltd., Kendall Square Research, Encore, Parsys, and Thinking Machines. In addition, we have announced plans to adapt our software to massively parallel supercomputers from Pyramid and Unisys(Intel) Corporations within the next year, and we are negotiating with additional vendors to support other massively parallel machines, which are as yet unannounced.

Historically, Oracle's revenue has derived mainly from servers in the midrange of computer power, between desktop systems and mainframes. Although our software products run on mainframes, we have not experienced the revenue growth in mainframe software that we have seen in the midrange. We believe that there are two reasons for this: First, mainframes are an increasingly costly technology. Over the last ten years, mainframe technology has declined in cost by less than a factor of ten, while desktop technology--now shared by mid-range systems-- has declined by a factor of several hundred. Second, the complex and proprietary nature of the mainframe operating systems environments has discouraged customers from committing new applications to them.

At the same time, over the past few years particularly, the rate of generation of computer data has skyrocketed in the commercial world: Sources such as scanners, mail and telephone order merchandising both at the wholesale and retail level, automatic teller and ticket vending machines, and credit and debit card purchasing are combining to produce enormous volumes of data. This data is generated in transactions which are presented to on-line systems at prodigious rates--rates which tax the most advanced traditional commercial systems produced today. Further, it has been demonstrated repeatedly that this data contains a wealth of information which can provide significant value to the enterprise capable of interpreting and exploiting it. To the extent that mainframe systems have been brought to bear on such "decision support" or "data mining" applications, they have been forced to deal with statistical summaries. It has been shown repeatedly, however, that for many purposes, "God is in the details" and summaries do not suffice. More powerful computing capability is required.

It is Oracle's conviction that the solution to both the capture and the mining of this massive amount of data lies with massively parallel supercomputing technology. Why massively parallel supercomputers? The answer for capture lies with the broad "scalability" of massively parallel systems which will enable them to produce far higher transaction rates than those now possible. "Scalability" means that, unlike conventional architectures, massively parallel supercomputers can add computing power, data storage capacity, and input/output capability with near linear gains over a very broad range which stretches far beyond that of the largest mainframes. The answer for mining or decision support lies partly also with scalability, but in addition partly in the structure of our database and application software.

Parallelism is an intuitive concept, but writing parallel software is difficult today, and this presents a significant barrier to reaping the benefits of Massively Parallel Supercomputers. I believe this is as much a result of thirty years of creating software for serial computer architectures as of anything else, and that this situation will change over time as we develop paradigms for parallel programming as powerful as those we have created for serial software. The fact remains, however, that today the most successful applications on massively parallel architectures are those for which the intrinsic parallelism is fairly clear. Finite element computations and Monte Carlo simulations are two examples from the numerical world of science and engineering. Relational database systems are another, from the world of non-numeric commercial computing.

A relational database system represents data in the intuitively appealing form of tables; rows and columns of these tables are manipulated by searching, sorting, combining, separating, partitioning, etc. to generate other tables containing the desired results. There is an international standard programming language, SQL (Structured Query Language), for carrying out these manipulations. Many, though not all, operations are intrinsically parallel, and it is this intrinsic parallelism which lends itself to exploitation by massively parallel supercomputers. That said, the creation of the software to exploit massively parallel systems is far from straightforward. Oracle has been adapting its software to exploit these architectures for over five years, at a cost to date of over $11M. We have made this investment because we are committed to this technology as the key to our growth at the high end of the market. We have a significant lead over our competitors, but even with the major investments we have made, we are not yet satisfied. We estimate that we will require another twelve to eighteen months to bring the software to a point at which its ability to exploit massively parallel supercomputers will be comparable, relatively speaking, to its ability to exploit the then-current mid-range or mainframe computers. Nevertheless, we are determined to make the additional investment because everything we have seen to date reinforces our conviction that we are on the right track.

One need only look at the list of vendors above and compare it to the list of a year ago to see that we are not alone. Several of our competitors have also indicated by their announcements and actions that they share this vision. As further validation of this view, I would like to share with you a quote from a study by International Data Corporation entitled: "Massively Parallel Computing: Opportunities in the Marketplace", which Oracle consented to join with a number of other companies in sponsoring earlier this year. IDC is a well-known and authoritative independent market research company in the computer industry. The IDC study surveyed a total of over 200 enterprises across a broad range of industries; all were Fortune 500 companies in size.

"Massively parallel processing is an idea whose time has come. The real question is whether the technology can match the idea. Over 16% of the surveyed population has already investigated MPP as an option for commercial applications....MPP's potential is significant across most of these industrial groupings, but, from their perspective, *the technology must first become a product*[sic]– accompanied by the service, support and reputation to carry it to success. This is not unlike the situation faced by most new technologies that are complex and hard to measure form the end users' perspective...".

These results were obtained from a survey population with these average characteristics:

- A large mainframe, associated with a large and rapidly growing database;
- A need for real time or near-real-time decision support systems;
- High volumes of user network traffic;
- Perception that data and decision models offer strategic competitive advantages;
- A need for high performance commercial systems which are flexibly configured to meet almost constantly changing information needs and flows.

We believe this survey encompasses some of the most visionary and competitive of our country's enterprises–the leaders, whom others will soon follow. It is particularly noteworthy that high volumes of network traffic go hand-in-hand with the perceived need for massively parallel supercomputing. I believe that this is a pattern which will persist, and will accelerate with deployment of the NII.

I would like to emphasize that relational database is but one area in which Oracle sees its future tied to massive parallelism. We have a text search and retrieval product which is also being ported to massively parallel supercomputers. It will depend on massively parallel platforms to enable it to process huge amounts of textual data, such as that in the Patent and Trademark Office, with which we are working on a joint project. We also have an Office Systems product, which offers electronic mail, meeting scheduling, and a variety of other work group and enterprise office functions. That Office software is running today in the production network at Oracle headquarters on a massively parallel system from nCUBE. By the end of 1993, our plan is to have nearly a thousand of our employees depending on this system for all of their office software functions. I should add that within Oracle everything from ordering office supplies to reserving conference rooms depends on electronic mail; it is central to our operation. Finally, we have announced an agreement with a Regional Bell Operating Company to perform a trial of electronic service to the home, including messaging, text retrieval and video-on-demand, all to be based on software running on massively parallel supercomputers. These supercomputers will not only provide the services, but gather the data required for accounting, billing, and other administrative operations.

That, Mr. Chairman and Members of the committee, is a brief snapshot of Oracle's commitment to massively parallel supercomputer technology. It should be obvious that we regard it as a crucial element of our continued success. This is a very different picture of the commercial potential for the technology than that painted by the CBO study. This committee and the HPCC are in my view to be strongly commended for your early support of massively parallel technology.

As requested, I would now like to focus briefly on the future of HPCC. There are some important areas which I feel that Congress and HPCC should consider as targets for its funding. These opinions are my own, and do not necessarily reflect an Oracle position.

Funding Software Development

I believe that the most important areas for funding now are those of software. I have a different viewpoint from that of most of those who have appeared before you, however, because I speak from the perspective of non-numeric commercial software. One major difference between my kind of commercial software environment and that of, say, an oil company 's geophysical computations or my co-panelist's financial calculations is in the duration of the application. Most of the calculations performed on supercomputers historically have been closed-ended; that is, they have a beginning and an end. At the end of one set of calculations, before the supercomputer is dedicated to the next, it can be re-initialized, service scheduled, software upgraded, and so on. That is much less true in the environment of most commercial applications for global enterprises. Electronic mail, credit processing, electronic services to the home, and many, many others are 5-7 day per week, 18-24 hour per day operations, and they demand new software design and functionality for massively parallel supercomputers. I will list five areas in which I believe there is a significant need for research work.

1. The need for Continuous Operation
First is continuous operation. It must become possible to operate these systems for many months at a time without interruption for any reason. To begin with, this means that both hardware and software failures must be "soft", i.e., capable of being isolated to the smallest possible set of components, without disruption to the rest. In addition, it must become possible for the surviving components to assume the processing burden of the failed ones in a rapid and graceful way, again without disruption to the processing which was untouched by the failure. Next, it must become possible to change the configuration of the system during operation. It must be possible to add capacity, to perform routine maintenance, and to upgrade the software processor-by-processor, for example, all without disruption. The US telephone system, as we know it as users, is a good paradigm for the reliability we must attain; it isn't failure proof, but its failures are generally few, and only very rarely are they disruptive in a significant way. This is definitely not a problem which can be solved with redundant computer hardware; to get there from where we are today on massively parallel supercomputers will require significant innovation in software.

2. The Need for Software to Manage Massively Parallel Systems
Of at least equal importance is work on software to manage massively parallel systems in operation. In addition to the "grace under pressure" requirements above, we need new techniques for managing large numbers of homogeneous, closely cooperating software systems: starting them up, pausing or stopping them, measuring what they are doing and how they are interacting with one another, and so on. In some cases, we are not sure today what are the best parameters to measure, in others we're not sure what is the best way to control operation. Finally, we need to understand how best to present the data to a human operator for most effective control.

3. The Need for More Advanced Operating Systems
If massively parallel systems are to undergo a level of refinement analogous to that experienced by conventional architecture, additional work is needed in the area of operating systems. New methods for communication and coordination among processors' software, new methods for load balancing and scheduling across processors, and ways in which to allow multiclass priority scheduling of shared disk operations are all fruitful areas for innovation, among others. Oracle is currently participating in a CRADA with Sandia National Laboratories and nCUBE for work on operating systems specifically targeted to support database software execution on massively parallel machines.

4. The Need for Additional Commercial Application Systems
Fourth, funding is needed for additional commercial application subsystems. Oracle can be successful in exploiting massively parallel supercomputer systems because we can hide the parallelism from a population of software designers who need only continue to think and write in their accustomed way. What they write, SQL (Structures Query Language), is a programming language which specifies tasks which are intrinsically parallel, but it is the Oracle database software which ferrets out the parallelism and exploits the hardware. We assume all of the burden of complexity inherent in writing the software which does this.

I believe that there are other analogous application areas, i. e., where the average application program can be written in a language which is intuitive to the serially-trained software mind, but the implied parallelism can be exposed and exploited by a small team of highly skilled parallel experts. This is certainly true in the design of large-scale logic chips, where the VHDL language specifies operations which are intrinsically parallel. I believe that it is also true in the area of "Expert Systems", where the underlying algorithmic technique is that of nondeterministic parallel search. I am convinced that there are also other areas, unknown to me, or perhaps even as yet unconceived by anyone, as were spreadsheets in the early days of personal computers. I encourage HPCC to look for such application areas for funding.

5. The Need for New Programming Tools and Techniques
Finally, I indicated earlier that I believe that we will, over time, develop paradigms for parallel programming which will make it much more accessible to the general population of software designers. This will not happen without an active search, however. As an example, I understand that investigations are underway in some locations on software architectures which isolate complex functions to specific and highly localized areas of a large software design. One might conceive that such an approach could isolate the areas where complex parallel programming is required. Of course, these complex areas must account for the bulk of the execution time of the software in order for parallelism to contribute in a substantial way to performance. I strongly urge HPCC to seek projects which will expedite the process of making parallel power accessible to the general population of software designers.

Supercomputer Hardware Issues

I have devoted all of my testimony up to this point to software issues, in keeping with the role in which you have invited me here. However, I have spent much of my career in the development of computer hardware, including both vector and massively parallel systems, and I feel compelled to comment on the naive notion advanced in the CBO report that floating point capability alone distinguishes a supercomputer from one which is not. The art of design of arithmetic units for both scalar and vector floating point computation is in an advanced state. While I do not wish to demean the creative ideas which advance this state, it has been my experience that other factors are today much more significant in the overall performance of supercomputers on real-world computations--as distinct from industry benchmarks--than the details of arithmetic unit design.

Memory design, input/output capability, processor performance on operating system specific tasks such as interrupts, and other factors contribute in major ways to the overall practical performance of supercomputers, or other computers for that matter. One major foreign manufacturer of vector machines learned about some of these factors the hard way a few years back, by failing to give them sufficient attention. In the case of massively parallel supercomputers, scalability of these and other factors is an added crucial element of the design. From Oracle's perspective, we need supercomputing capability in all of these other factors just as the floating point intensive users do, even though floating point performance is largely incidental to our current applications. When evaluating hardware proposals, I strongly urge HPCC to look beyond mere floating point (benchmark) performance, contrary to the viewpoint taken in the CBO report.

Mr. Chairman and Members of the committee, thank you for the opportunity to present a view of supercomputing technology supported by HPCC as seen through the eyes of one long actively engaged in the development of commercial software. I hope that my comments will be helpful in enabling you to put the advances in massively parallel supercomputing fostered in part by the support of this committee and HPCC in better perspective. As a beneficiary of past funding for massively parallel technology, we at Oracle hope that you will continue to support the next phase of research need in software for such systems

W. Donald Frazer is Vice-President, Massively Parallel Products at Oracle Corporation, the world's largest supplier of data management software and services.

He is responsible for Oracle's products which run on massively parallel hardware platforms. These platforms include those of all current suppliers to the commercial marketplace. He has been a well-known figure in the computer industry for over thirty years, with a career which has encompassed industrial research, academia, and both hardware and software development. His long association with high performance computing has included responsibility for IBM's first vector processor products, and massively parallel systems at Teradata Corporation, in addition to his current role.

Dr. Frazer holds a BSE from Princeton University, and MS and PhD degrees from the University of Illinois, all in Electrical Engineering.

Mr. BOUCHER. Thank you very much, Dr. Frazer.

Dr. Bloomquist.

Dr. BLOOMQUIST. Thank you, Mr. Chairman. I am here representing Mobil Oil, a Virginia-based company, I might point out. We have a massively parallel computer in operation today which I recently was involved in convincing my management to spend several million dollars for, and I assure you we're out to make this computer do things that are useful after spending that kind of money.

But that system today is operational in Dallas, Texas, providing some high performance, high technology jobs in Dallas and hope to continue to do so. We bought it from a company based in Boston, Massachusetts, and the computer itself is made of much of American-made components. High performance computing have been an area of leadership in the United States for many years, and I hope will continue to be so.

We're not alone in the petroleum industry. In fact, the petroleum industry is one of the largest users of high performance computing, and I was somewhat surprised to notice that the Congressional Budget Office report did not make much of that, but we use about 20 percent of the private sector high performance computing cycles today, and have for many years. There are a number of other petroleum companies who have made investments in massively parallel computers. At least two other major U.S. oil companies and several service companies have recently made million dollar or more purchases of massively parallel systems.

The thing about these systems is that we're using them to do things that we couldn't do before. I think some of the previous speakers have noted that that is really the true benefit of these types of systems. We're able to, for example, image complex geologic structures in places like the overthrust region in the western U.S. or below the salt in the Gulf of Mexico that we haven't been able to image before. And certainly there's a potential for finding some major new oil reserves in places like this that could certainly improve our balance of payments by having to import less oil and generate more jobs for Americans in the service sector of the petroleum industry to develop those new oil fields, if any are found.

We're also using them to better model how existing oil fields can be produced and optimize that production to squeeze more oil out of the ground and to do it with the most cost-effective investments. And we've used it also for some fundamental research. In fact, one of the programs that we've done has been with Los Alamos research facility. We did some fundamental modeling of fluid flow in the earth that helps us better understand how to produce oil and gas; some very fundamental understanding of the wetability of rocks that makes a big difference in how much oil you can get out of the ground.

I also would like to say in terms of the program here that we think it is important that you have a balanced program, but the massively parallel computer is certainly the wave of the future and it is the correct area to be emphasizing. But in addition to very high speed processors, you do have to have the software that goes with it. You do have to have high communication networks to get data in and out of the machine and to get them to the people who are ultimately interpreting the data, and you have to have high performance workstations. You have to have high visualization systems and high performance graphics to be able to analyze the huge amounts of data that you can generate with these systems. So it is important to look at the overall picture and not just concentrate on building a faster processor, but instead to look at the entire process and the system that we're going to use these systems for.

Finally, in closing, it was interesting, yesterday when I was going over the testimony that I was intending to give today with Maury Devine in our Government Relations Office, Maury's not a technical person, but after I told her what we're really talking about here today, she said: "You know what this really is all about in many respects is innovation and creativity and what we're really trying to do here is give our scientists the tools to be innovative and creative, much as they've been in the past." I think American science takes a back seat to no one, but by providing the kind of tools that we can give them with high performance computing, I really think we'll be able to maintain that kind of leadership, and I look forward to being part of it in the future.

Thank you.

[The prepared statement of Dr. Bloomquist follows:]

Testimony on High Performance Computing and Communications

Before the House Committee on Science, Space, and Technology

Subcommittee on Science

by Marvin G. Bloomquist

Manager, Information Technology

Mobil Exploration and Producing Technical Center

October 26, 1993

Testimony on High Performance Computing and Communications Before the House Committee on Science, Space, and Technology
Subcommittee on Science

by Marvin G. Bloomquist
Manager, Information Technology
Mobil Exploration and Producing Technical Center
October 26, 1993

Mr. Chairman and Members of the Subcommittee:

I am Marvin Bloomquist, Manager of Information Technology at Mobil Oil Corporation's Exploration and Producing Technical Center in Dallas, Texas. Our organization is a primary technology support center to Mobil's upstream exploration and producing activities. High performance computing and communications are major factors in our support of the company's core business objectives of finding, developing, and producing hydrocarbons. The petroleum industry has been one of the largest private sector users of high performance computing for more than 20 years.

I'm pleased to be here today in support of the High Performance Computing and Communications (HPCC) Program. However, I encourage you to take a balanced approach to complete the HPCC Program. We are especially concerned that an equitable level of investment be included for massively parallel supercomputer systems and software. It is balance that we are looking for in the program and this is consistent with Mobil's policy to encourage a variety of solutions to technological problems. The petroleum industry is just beginning to make significant investments in massively parallel technology. I believe it will play an important role in our industry in the near future.

High Performance computing is an important tool in the petroleum industry and a critical factor in the competitiveness of many segments of American industry. Tangible results from the HPCC Program are being utilized in the petroleum industry today. I would like to share with you how Mobil has been able to capitalize on the HPCC initiatives in applying massively parallel supercomputing technology to seismic imaging of the earth.

The History of Supercomputing at Mobil

Mobil has a long history of involvement with supercomputers. In the mid-60's Mobil and several other oil companies sponsored research at the Massachusetts Institute of Technology to develop basic computer algorithms needed to process seismic data. We concluded that this new digital seismic method was a major technical advance, improving our ability to identify favorable geological structures. In 1967, Mobil acquired a Control Data Corporation Model 6600 to exploit this technology.

This was the first "supercomputer" to be purchased by an oil company, and it was key to establishing Mobil as an early leader in digital seismic technology. The late 60's and early 70's were very successful years for finding new hydrocarbon resources, aided by this new seismic technology.

In April, 1989, the first massively parallel supercomputer in the petroleum industry was installed at Mobil's Dallas Research Laboratory. Thinking Machines Corporation formed a partnership with us to establish the viability of this new technology in the petroleum industry.

By the end of 1989, significant progress had been made in implementation of key portions of seismic and reservoir modeling processes on the massively parallel supercomputer. For one of these processes, Mobil and Thinking Machines jointly received the Institute of Electrical and Electronic Engineers' Gordon Bell Award for the fastest practical computer program in the world.

In 1990, the research program was broadened and the hardware was upgraded. A larger research and development team was formed to translate these experimental results into a practical seismic data processing system. This required both the development of more sophisticated parallel software and high speed data handling systems to feed the machine at a sufficient rate to take advantage of the computer's processing speed.

In February, 1992, we installed an early model of Thinking Machines' newest model, the CM-5. Development of additional seismic applications continued as the system matured. At the same time, the project team had been integrating other elements of the supercomputing environment: workstations, high-density mass storage, and high-speed communications links. We benefited from pioneering work from the HPCC program in several of these areas including development done at Los Alamos National Laboratory, University of Illinois, and University of Pittsburgh.

Business and Technology Drivers

Obtaining good quality seismic images for reservoir development and exploration is often complicated by unfavorable subsurface acoustic conditions. Under these conditions, conventional seismic technology may incorrectly position geological features, or not image them at all. While more accurate methods are known, they require weeks or even years of computing time on conventional vector supercomputers. The potential of these more accurate processing methods have been demonstrated on a limited scale in experimental work documented in the literature by Mobil and others.

These improved seismic imaging methods could open up new exploration plays in extremely complicated geological conditions such as exist in the overthrust region of the Western United States and below layers of salt in the Gulf of Mexico. A major discovery in these areas could decrease the United State's dependence on imported oil and also create jobs in the oil service industries. In existing oil fields, better seismic images of complicated subsurface geology can reduce development costs and increase the amount of hydrocarbons recovered.

Throughout the petroleum industry, efforts are underway to commercially develop one of these more accurate seismic methods called prestack depth migration. This is an extremely compute intensive process. Several companies have demonstrated this technology in "target-oriented" applications where a small portion of data from a large field seismic survey is used to image only a specific area in the subsurface. A much greater degree of accuracy could be achieved if the entire data set covering the area could be imaged. However, such "full-volume" imaging using prestack depth migration algorithms requires trillions of calculations and tens of gigaflops of computing power. The only way this amount of computing power can be delivered today at an economically viable cost is with a massively parallel computer. Larger surveys could easily utilize teraflop size machines if they were available today at reasonable cost.

As an example, an exploration seismic processing problem representing the subsurface beneath a 25 square kilometer area on the earth's surface can be completed in weeks on a 64 processor massively parallel supercomputer instead of months on a conventional vector supercomputer. The cost to do this processing with a vector computer would be prohibitively expensive. The problem is clearly not feasible with workstations today, even if multiple workstations can be efficiently used in parallel. Contractors providing seismic processing services to the petroleum industry believe within the next two years, prestack depth migration in three-dimensions may become a commercially-viable service using massively parallel computers.

During the past two years Mobil, at least two other oil companies, and several seismic processing service companies and have purchased massively parallel computers costing several million dollars each. In today's cost conscious environment, I assure you that this level of expenditure would not have been made by any of these companies with out a reasonable chance of being able to gain economic benefit from the investment.

Future Development

Despite the success achieved in applying massively parallel computers to seismic imaging, the potential of this technology has just begun to be tapped. To take advantage of parallel processing frequently requires rethinking the problem solution.

Our mind has been taught to solve problems sequentially for so long, it has to be retrained to solve problems in parallel. Massively parallel computers need to be available to the academic community, not only to solve new problems, but also in order to develop the skills among the scientific and engineering community needed to utilize these machines. This type of funding should be a high priority for the HPCC.

Successful application of high performance computing usually requires a balance between processors, communications channels, storage media, and graphic display systems. Without attention to all these, a computing system cannot achieve true high performance. HPCC needs to properly balance support of all these areas.

Conclusions and Recommendations

In summation, massively parallel architectures are the supercomputers of the future, but we are able to use them to do useful work now. These systems are being applied to the solution of problems which cannot be accommodated by conventional supercomputers.

The transition from one computer architecture to another is a complex and time-consuming process. The HPCC Program has enabled the hardware transition.

The larger challenge now becomes the evolution of the software environment required to enhance the productivity potential offered by the new architecture.

Historically, advancements in computer technology have come about as a result of cooperative efforts involving research institutions and industry, with government sponsorship. This is the case today, and I recommend continued, balanced appropriations to bring the program to its planned conclusion.

The benefit of the HPCC Program to the petroleum industry is a pronounced enhancement of our ability to meet our most important business objectives - finding and producing hydrocarbons. The American people will benefit by the technological advantages that leadership in high performance computing provides to industry and the academic community. We will all benefit by retaining high-paying technology jobs in the United States and giving American industry a competitive edge through technology.

####

Marvin G. Bloomquist
Manager, Information Technology
Mobil Exploration and Producing Technical Center
Mobil Research and Development Corporation

Marvin Bloomquist is Manager of Information Technology at Mobil Exploration and Producing Technical Center in Dallas, Texas. He has been with Mobil since 1968 in a variety of positions, mostly associated with geophysical research and development. Since 1988, he has been closely associated with the application of massively parallel computers.
Mr. Bloomquist is a 1967 graduate of the University of Texas with a Ph.D. in Electrical Engineering.

Mr. BOUCHER. Thank you very much, Dr. Bloomquist.

We want to express our appreciation to each of the witnesses who have been here this morning for their very useful comments. The written testimony will be particularly beneficial to us as we work with the various agencies that are responsible for carrying out the high performance computing program and fine tuning their mission over time and giving them somewhat better guidance.

ADVISORY COMMITTEE

I'd like to ask questions in a couple of areas. One of the shortcomings to date of the program is the fact that the High Performance Computing Advisory Committee, which was mandated in the 1991 legislation as several witnesses have indicated, has not been appointed as of this date. That committee was designed to be the principal mechanism for receipt by the federal agencies of advice and suggestions from the private sector and from the users of this research as to how the research could best be targeted, how allocations could be made, and the kinds of research projects that ought to be carried forward. That still hasn't happened.

I'd like to get your thoughts about just how important it is that we encourage the Administration to appoint that committee at the earliest possible time. I have some reason to believe that the Administration has intent to do that, but it certainly wouldn't hurt to have a distinguished panel of witnesses such as this endorse that approach and tell us why that's important. So I guess that's what we lawyers would call a leading question.

If anyone would like to begin—Dr. Rubbert, we'd like to hear from you.

Dr. RUBBERT. Yes, I would like to comment on that. I think that it is critically important that we take the steps that will lead to industry having an active involvement in the planning process, not only the downstream execution. This is a project in which industry is the customer. We're trying to improve the competitiveness of industry by influencing our design processes, and so forth, and in order to be effective in doing that, it is absolutely imperative, in my opinion, that industry be involved upfront in the initial planning of the program.

Thank you.

Mr. BOUCHER. Thank you very much.

Dr. Berlin.

Dr. BERLIN. Yes, Mr. Chairman, I think it's such a serious situation that if the—if this continues and the Advisory Panel is not put in place and there's not some mechanism for linking user needs to the program, the program could literally float away. And if we look back, when the program was originally set out, the Science Advisor, Mr. Reagan's Science Advisor, specifically dipped into the agencies and pulled out users, and those users spent about six months analyzing the state of the technology and the state of the applications, and they set the first road map for the HPCC initiative. It has been in the last three years that user insight has fundamentally been lost and there is no mechanism whatsoever by which it is possible even for the National Coordination Office—they can't meet with industry to plan because they get in trouble with FCCA and things like that.

So I would specifically recommend for your consideration, Mr. Chairman, that when H.R. 1757 goes to conference, that you might consider working with your colleagues in amending it to require the President to, within 30 days, get that advisory board done.

Another possibility might be to go ahead and lower it to Dr. Gibbons' level, so they could just appoint it, and it doesn't go through presidential personnel.

Mr. BOUCHER. Dr. Berlin, let me be devil's advocate with you for just a moment.

Dr. BERLIN. Okay.

Mr. BOUCHER. You said there's no mechanism by which the private sector can interact with those responsible for——

Dr. BERLIN. That's right.

Mr. BOUCHER [continuing]. Coordinating the HPCC Program. Do not the various agencies, however, that are responsible for carrying out the direct research mission have in the main external advisory bodies that perform that function, and have not those bodies had some ability to influence the course of the HPCC Program as carried out by those agencies?

Dr. BERLIN. They do have advisory boards, but, by and large, the second question, the answer is no. If you look at the Defense Science Board, for example, which is one of the most active of the advisory bodies, the Defense Science Board has done some work with the services, but it has never been called in to assess ARPA's role, and that, of course, is the largest expenditure in the HPC program.

To my knowledge, there has never been any assessment of direction or any advice sought. There's been a lot of advice volunteered by individuals, but it's pretty much a one-way communication.

Mr. BOUCHER. All right, yes, let's hear from Dr. Bridenbaugh.

Dr. BRIDENBAUGH. I'd like to underscore the need for getting a balanced view from both industry and academia. All users are not created equal, and the applications, for example, that we would—that would be of interest to people in the materials business might be substantially different than those in the chemical or oil business, and the interest of academic researchers may or may not coincide with the needs of industry. So I think it's critical that you have in place a mechanism that provides a balanced view of both academia and industry, and not just people that are in the computing business in industry, but also user groups like Alcoa and Mobil and the car companies.

Thank you.

Mr. BOUCHER. Okay, Dr. Frazer?

Dr. FRAZER. I'd like to add my support to the suggestion for an overall industry input. One of the things that seems to me such a group could provide is guidance on themes that go across various agencies which may have more parochial interest in this application area or another. And one of my pets, which I think is really important, is the creation of tools and techniques to make the programming of massively parallel machines available to the vast population of application programmers. It's a very complicated discipline, and I believe that there are ways—and I've made a couple of suggestions in my written testimony—for research areas that can really broaden the scope of people who are able to capitalize on this technology and enable it to reach its full potential.

Mr. BOUCHER. All right, any further comments on this question? [No response.]

RESOURCE ALLOCATION

Let me move on to another issue, and that is the question of allocation of resources among the various program components. As of the present time, the high performance computing systems have received 22 percent of the funding. The advanced software technology has received 43 percent. The National Research and Education Network has received 15 percent, and basic research and human resources have received 20 percent of the funding.

Very briefly, if each of you or those of you who choose to do so, would you comment on the appropriateness of those allocations? And if you have suggestions for how they should be different, please tell us that and give us the supporting rationale that you have for why the allocations ought to be different from that.

Any volunteers? Dr. Rubbert.

Dr. RUBBERT. I would like to make a comment for the aerospace industry, and I don't really mean that it spill over into the other industries because I really can't talk for those industries.

But in our industry we have found that the companies themselves have to write the application software. We use from academia and NASA research laboratories and so forth, the intellectual knowledge, the technology for how to do it, but the actual implementation in terms of application software is done by the aerospace industry.

And my view of massively parallel today is it is not quite there yet for our industry. It may be for others, and we've heard some testimony that indicates it is. And so I view in the near future for my industry at least the role of massively parallel application software is really to exercise the hardware, find out what it's good at and not good at, and use that as a means to get the hardware developed to the point where it is really the best thing around and then build the application software. So we have to be careful in timing of when we do our investments in application software.

Mr. BOUCHER. Do you have any general comment about this range of allocation, however? I think your comment about how we deal with massively parallel computing software could be dealt with within the general allocations we have here. Are they basically proper as far as you're concerned?

Dr. RUBBERT. I don't know.

Mr. BOUCHER. No? Okay.

Dr. Frazer, would you like to comment on that?

Dr. FRAZER. Just quickly. You know my bias from my testimony. It's very much toward software.

It seems to me, given the menu of current and imminent suppliers of massively parallel hardware platforms, that the natural forces of competitiveness are at work and advances in the hardware technology are going to come fairly automatically simply because people are going to be competing for the same market. So I would suspect that you could shift the allocation from hardware, some of the allocation from hardware more into software. I'm not prepared to comment on the——

Mr. BOUCHER. Software already has 43 percent of the total program. That's by far the largest program component. Are you arguing that it ought to be higher still?

Dr. FRAZER. I would. Typically, in a hardware manufacturer, you will find a hardware manufacturer that supplies an operating system and supporting software, that there are more than twice as many—more than twice as much is invested in the software as in the hardware development.

Mr. BOUCHER. Does anyone have a differing view from that or just an alternative view? Dr. Audley.

Dr. AUDLEY. I would just observe that I'm surprised at the allocation. I would have thought that, based on the culture that I hear, that the hardware allocation was higher. I would also observe that the—and I'm delighted with the 43 percent devoted to software. I'm also delighted at the 20 percent devoted to the basic areas because I feel as though this is a long need item area that eventually feeds into the software and into the other areas and might, since it is also an intellectual enterprise, be considered as software as well.

Mr. BOUCHER. All right. Yes, Mr.—Dr. Bloomquist?

Dr. BLOOMQUIST. It just seems to me that the allocation is about right. In fact, if you drain off additional funds to put into software development, the other areas are already at a fairly low level of funding and it does seem to me to be a balance. These are big machines. It costs a lot of money to build, and you do have to put some money in hardware.

Mr. BOUCHER. Okay. Dr. Berlin, do you care to comment?

Dr. BERLIN. Just one comment and that's referring back to our discussion about the advisory panel in this context. This is a good example of why we need an advisory panel. The allocation is actually very deceptive because where the real question is answered is in what that means and which projects specifically are being funded, and what are the directions. There is no visibility in that.

What it does is it tends to reduce the number of constituents for the program out in industry and the number of people who see themselves as potentially being involved in what is a very, very important national program.

Mr. BOUCHER. I think you're exactly right about that. This subcommittee is very poorly positioned to try to make decisions or even make recommendations in terms of the kinds of allocations that ought to take place. That is precisely the function of a coordinating committee comprised primarily of the user community, and that's the role we would see that committee performing in the main with a lot of subsets.

What I was trying to do was just get a sense of whether or not we're badly astray in the absence of the committee having been appointed to make those recommendations, and it appears that we're probably not that badly astray.

I have a number of other questions, but before turning to those I'd like to give other members of the subcommittee an opportunity, and I would call on Mr. Smith first.

Mr. SMITH. Thank you, Mr. Chairman.

One message I think I hear is the importance of the HPCC to industry. Let me carry that over and get your justification or feeling why federal appropriations are necessary.

I've got two questions. One, why are federal funds necessary as you see it? And shouldn't industry be doing this themselves and coordinating this kind of effort within industry? Then my second question is going to be on the competitiveness competition factor. So maybe on the first question you could respond—whoever. Yes, Dr. Bloomquist?

Dr. BLOOMQUIST. Well, just—I consider us to be somewhat on the leading edge of this technology at this time, and certainly we have benefited from some of the work done with the HPCC. I don't think we could be as far along as we've been today—the University of Pittsburgh and Illinois and Los Alamos all, for example, pioneered some of the high performance communications links. We need to be able to get data in one of these machines to use it effectively for seismic processing. Without them having done that, I don't believe that today we would be using these machines in a useful way.

Certainly some of the software development that's going on as well, we're not going to program these machines in machine language. We don't have the resources to do that in our company, and we're one of the biggest oil companies in the world. If that kind of software isn't developed by some others, then we're not going to be able to use the machines, and that's another area that's certainly been funded by HPCC.

Mr. SMITH. But let me follow up on that still. If—are you suggesting that the demand is not great enough that the private sector would develop that software if Federal Government wasn't in there with whatever, 1.7 billion or something?

Dr. BLOOMQUIST. Well, we're certainly spending a lot of our own funds on developing software, but some of the more fundamental stuff we don't have the skill set within our company to do it or the resources to do that, and the risk and the benefits are not great enough for us without that—that work can be shared by many others; it's not strictly applicable to the oil industry. It's applicable in many other companies and many other industries as well.

Mr. SMITH. Dr. Rubbert, you had a comment? Yes?

Dr. RUBBERT. I'd like to comment on that question. HPCC is a very, very broad program, spanning both traditional areas of science as well as becoming involved in industrial processes. It is certainly proper and traditional for the Federal Government to be involved in the support of science, education, natural defense issues, et cetera, and HPCC does support those needs.

I think the real question is: how far down into the industrial world do you carry it also? We're living in a world economy where our trading partners do provide selective support for their industries, and so I feel that it is appropriate to take advantage of the HPCC technology to try and make it useful and effective for the U.S. industry.

I think we have to be a little bit careful on how we do that. The traditional way we do science is for a visionary laboratory, or whoever it is, to assume a strong central role as the customer and then people that help execute the program are suppliers. But when we're talking about helping industry do things better, we're really faced with a role reversal. Industry is the customer, and Federal laboratories, academia, and so forth, that support that are the suppliers.

So it's a culture change and you have to operate in a different way in terms of government involvement and laboratory support, depending on whether you're working on science or working on supporting engineering goals. And so I think it's appropriate. You've got to be careful how you do it.

Mr. SMITH. Let's follow it up with the competitiveness question.

Dr. RUBBERT. Yes.

Mr. SMITH. And I am becoming more selfish about applying our federal research money and our efforts such as HPCC that can be better utilized and more effectively utilized in this country to give us whatever competitive advantage that we can.

If all industries have it and if this information is public, and if the industry in other countries has as much access to it as we do, then I'm somewhat troubled by the competitive advantage that our money might otherwise give us as opposed to other countries. So help me understand that.

Dr. RUBBERT. I would like to speak to that one also. The name of the game in industrial competitiveness, as more and more we're coming to understand, is cycle time. And, yes, our European competitors for the aerospace industry—that's Airbus Industries—they can buy the same computers as we buy because our vendors sell them worldwide, but where we capture our competitive advantage is, by becoming involved in the HPCC Program and helping to execute it, we get a much better view of how the technology is coming along, which means we can make earlier and better and more informed decisions on what role it should play for us, what computers we buy, and so forth.

And the fact that we have been involved in the development of the software as well means we can get there first with software development and make it effective for us before our European competitors. Now that's the competitive advantage: to do the right things sooner.

Mr. SMITH. So—and just one last question, Mr. Chairman. What do you see as—the competitive advantages lead time, what normally, typically is going to be that lead time period? Are we going to have six months or one year because we're instrumental in the development of either the hardware or the software? What—give me a time frame on the lead time that's going to be essentially the competitive advantage for us?

Dr. RUBBERT. In my industry I would measure it in years rather than months. I really can't put a finger on whether we're talking about two years or three and a half or what it is.

Mr. SMITH. Dr. Frazer—Doctor?

Dr. FRAZER. Just a couple of comments. First, I think that you need to put it in context, historical context. Taking MPP as an example, HPCC was one of the very early sponsors of MPP when it was regarded as little more than a technical curiosity by a large percentage of the technical population both in industry and in universities. And it's through the sustained support of that over a period of several years that the technology was able to mature to a point at which its full potential could be understood and realized, and we're now on the threshold, we believe, of widespread exploitation of the technology.

So I think that one of the dimensions in evaluating government participation is the sponsorship of leading-edge technologies at an early stage to get a thorough understanding of their full potential, as has been the case with MPP.

I think that there is a role for government, and I would second the comments made on what's happening abroad. I'll point out that two weeks ago ICL Computers, Limited, which is a British company owned by Fujitsu, announced a massively parallel machine called Gold Rush which derived from a joint funding by the EC in which Bull and Olivetti, I think, were the other participants. So other countries or consortia of countries are, in fact, doing it, and it has come to be the way of the nineties, and I think we should do it as well.

Mr. SMITH. Mr. Chairman, thank you.

There was an—Dr. Bloomquist for a comment.

Dr. BLOOMQUIST. Just one very short remark: one of the problems with all this stuff is there's a knowledge base required to use it, and that's not something you can develop overnight or buy off the shelf. That kind of investment that does take years to build up the engineers and scientists with the knowledge to be able to use this type of equipment is something that I think we do have a competitive advantage over other countries that are investing in this technology today.

Mr. SMITH. Well, and likewise, some businesses in this country have the advantages over other businesses. So the superknowledge with supercompanies becomes a dominant force that's pushing aside other countries that are trying to develop it.

Mr. Chairman, my quandary—and maybe you can help me solve it sometime—is simply, as we've depended on the military justification for so much of our technological and scientific investment, and now as we pull away from the military, how much do we replace with other efforts to justify it?

Mr. BOUCHER. Thank you very much, Mr. Smith. That's a thoughtful question, and we certainly have a distinguished panel of witnesses here who can help contribute to that answer.

We will recognize subcommittee members in the order in which they arrived this morning. Mr. Minge, do you have questions for our witnesses?

Mr. MINGE. Yes, I have a couple of questions.

First, Dr. Ruppert—or Rubbert, sorry—I am very interested in your comments about the massively parallel computer systems and perhaps the overemphasis that we've had on those to the detriment of the development of other types of computer systems. Could you just explain briefly why you feel that the other systems have been neglected and what we should do in order to rectify that shortcoming or that slight?

Dr. RUBBERT. Well, I guess my point is the following: as we move forward with the massively parallel, we have to keep asking the question: how good are they? But you have to have a baseline against which to ask that question. The baseline are the nonparallel machines.

And we will take advantage of and use massively parallel effectively if, in fact, they do the overall job better than the other machines. And so we have to keep benchmarking against what you can do there.

I also conducted basically a technical assessment with my technical staff last week addressed at this specific question. Their present vision is that the type of architecture that seems most suited to our problems over the next five to ten years is probably the modestly parallel, 10, 20, 30 processors rather than massively, and I think it's driven by the particular character of the problems we're trying to solve.

Mr. MINGE. Do you feel that the Federal Government at this point has devoted perhaps excess resources to the development of massively parallel computers as compared to the other types of computers?

Dr. RUBBERT. Yes, I do. I think we need a little better balance. My exposure to HPCC indicates it's almost exclusively massively parallel. I think there needs to be a safety net of the nonmassively parallel computers and try to bring that along as well.

Mr. MINGE. Secondly, I'd like to pose a question to Dr. Rubbert—I'm sorry, Dr. Bridenbaugh. It's hard for me to keep all the names straight with the faces.

I notice that you're using a National Science Foundation computing center for your work at Alcoa, and I have heard some concern that when industry is using these centers it sometimes is receiving a Federal subsidy in a private sector that places other providers of computer services that are operating on a proprietary basis—that is, a business basis as opposed to a university center basis—at a disadvantage.

And I'm wondering if you could sort of speak to the question: when should the Federal Government be subsidizing these NSF centers for industrial use as compared to simply telling the private sector that you should buy computer time on the market?

Dr. BRIDENBAUGH. Well, first, let me comment that there is not a lot, that I'm aware, of privately-offered high performance computing time available on the market in shared services. There may well be and I'm just not aware of them.

And if—it's interesting that you phrased it the way you did because I thought that you might comment that industry is subsidizing the centers in many ways bec use we pay for the time and nobody else does. I mean, the universities, academics, they get free time through the NSF program.

And maybe I could also address this issue of lead time in the answer to your question. When we have looked at the problems of the kind that I talked—and maybe we'll talk about the aluminum-intensive vehicle—we have solved those problems through use of the Pittsburgh supercomputing center in one-tenth to one-twentieth of the time that it would have taken to do that on our own in-house machines. And that barrier was so large that we would never have attempted to solve the problem.

And then the consequence of that would have been that we wouldn't have designed and developed a car that we have designed and developed in this whole new concept without that capability or it would have taken us much longer. There's absolutely no way that we could afford to develop all of the high performance computing technology on our own. We can't afford, I think, to have our own in-house capability. We couldn't keep it running 24 hours a day seven days a week because we don't have many problems.

So access to those supercomputing centers has given us an opportunity to solve problems we never would have solved before, would never have gotten there without that technology being developed by somebody. And you could argue, of course, that if we could somehow magically get all of the companies in America to cooperate on any one given day and coordinate all of their activities, we could probably do this ourselves instead of having the Federal Government do it.

We have no mechanism to go around and tax all the individual—we have no mechanism to go around and tax all the individual companies and get them to the table. And so I think it's very appropriate use of Federal money to develop this technology, to make it accessible to industry. I think it's also very appropriate that industry pays to use the equipment and knowledge and software as well.

Mr. MINGE. Dr. Bloomquist, I have a question for you as well on—I believe you said that Mobil had purchased a computer. Was that within the last couple of years? How recently was that?

Dr. BLOOMQUIST. We recently purchased a larger massively parallel computer. We actually leased one earlier and have since 1989.

Mr. MINGE. I read with some interest the "Atlantic Monthly" article about the U.S. semiconductor industry and what happened in the 1980s, an article I think from November 1993. And my concern is the competitiveness of the American industry vis-a-vis the development of high performance computers in other countries that may have policies that protect or promote their developing industries more than we do.

And I'm wondering if in your purchase or leasing of a computer whether this question of the competitiveness of American industry versus, let's say, a Japanese-produced computer came in and what, if any, insights you have into what we should do to make sure that our high performance computing industry does not suffer like our semiconductor industry did.

Dr. BLOOMQUIST. Well, we certainly feel like we have to bring the best technology to bear regardless of where it comes from, and we evaluated supercomputers or massively parallel computers from a number of different manufacturers, including a Japanese manufacturer, before we made the purchase. I was pleased to see that the American computer that we did purchase we felt was the best piece of technology available.

I think much has been alluded to software here, and that's really the secret of these massively parallel computers. That is an area that I think we have excelled here in the U.S.—something about the American educational system or brain or something that seems to be well suited to software development. And if we can maintain that kind of lead in the software end of things, I think that's what really will drive the development of these computers, more so than hardware systems, though they certainly are an important part of it.

The hardware, and particularly the integrated circuit area, is, of course, where the Japanese have taken a lot of the technology from us, and they have the lead in some of those areas, but not in software. That's where U.S. companies still, I think, have an advantage.

Mr. MINGE. Thank you.

Mr. BOUCHER. Thank you very much, Mr. Minge.

Ms. Eshoo.

Ms. ESHOO. Thank you, Mr. Chairman, and welcome to all of you.

I'm always amazed that we don't have more members any time this subcommittee has hearings because they are the most instructive, and today does not step out of character. So thank you to each one of you; a special welcome to Dr. Frazer whose company, Oracle, is in my magnificent congressional district, and to Dr. Bridenbaugh, I—about 30 years ago this time, I worked for Alcoa. And so it's nice 30 years later to be working with you in a partnership here.

I think the best questions have already been asked. I think that, No. 1, this subcommittee needs to take a very proactive stance with the White House that the Advisory Committee be appointed forthwith. That has been established. It's needed. I would also like in that directive from the committee that women in significant roles played in America's corporations be included in that Advisory Committee. I'm always amazed that there aren't any that come forward to this witness table. So I have to put that plug in since we have some on the other side of the table here.

The question that I'd like to ask you is: in your view, since there is not an Advisory Committee, are you pleased with the outcome of the investments that we are making with the dollars in these four important areas? And how are they being measured? How is it being measured, or is it? Or is it just driven by these agencies without any kind of measurement?

Yes? Whomever.

Dr. RUBBERT. I would like to answer that from the aerospace industry. I am very pleased——

Ms. ESHOO. You're a great advocate for the aerospace industry.

Dr. RUBBERT [continuing]. With the outcome. However, there was wasted money and spun wheels. NASA, without talking to industry, put together a program plan and it was reviewed and, as a result of that review, we entered into a replanning process and did it again.

Ms. ESHOO. Yes.

Dr. RUBBERT. And so the outcome was good, but it wasn't the most efficient process.

Ms. ESHOO. Thank you.

Dr. BERLIN. Yes, I think that, by and large, those who have watched the program are pleased with the general outcome of where high performance computing has come over the last 10 years and its application in industry. However, I have to say that one of the problems is there is no one assessing the specifics for most of the programs.

And, by the way, Mr. Chairman, there was one advisory board that did meet and did this very excellent report. In fact, if it's not going to be in the record, I'd recommend that it be put in the record because it is the only document that really has ever effectively assessed any part of the HPC program and it did an excellent job. And Dr. Bridenbaugh was one of the users on that panel.

Mr. BOUCHER. Dr. Berlin, I appreciate that recommendation. We'll not print it as a part of the record itself, but we will receive it as part of our subcommittee general inventory of documents. It will be quite helpful for that purpose.

Dr. BERLIN. So I think that one of the problems is we don't have goals being set by users; we don't have users assessing where things are, and the sad part about it is that the people in the government program are very, very committed to helping America and helping their agencies do the job. These are not folks who are trying to hide under a bushel basket. They need the help of the external assessments, but it has to be done in a way that politically they can accept the results.

Ms. ESHOO. Is there anyone else that would like to comment?

Dr. AUDLEY. I would make one more comment on this subject. I think one of the things that the HPCC Program has been able to do is to provide the mechanism to inspire a new way of thinking in a lot of different areas.

Ms. ESHOO. That's remarkable for government, isn't it? [Laughter.]

Yes, it's so pleasant to hear, and coming from you that's wonderful.

Dr. AUDLEY. Well, I think ours is an industry that was very low technology prior to 1980. Since then, it has blossomed forth, I think, due in large part by these technologies that have begun to mature, and ours is an industry that hasn't necessarily benefited directly from government grants or other—for example, when we do our competitive searches, we don't use federal laboratories or federally-available resources, but we do find opportunities in the federal research budget that ultimately works its way into the commercial sector.

One of the things coming from the industry that I do, that I am concerned about, is that mechanisms be found whereby the federal research dollar somehow finds its soul brother in the venture capital arena to take forward a lot of the technologies that are germinated in government sponsorship.

Ms. ESHOO. Yes?

Dr. BRIDENBAUGH. Well, first, let me say I'm sorry you left our company 30 years ago. [Laughter.]

Ms. ESHOO. I made $350 a month; I'm not sorry. [Laughter.]

Dr. BRIDENBAUGH. Well, you would—you might still be making $350.

Ms. ESHOO. Of course, that was 30 years ago, but it was a wonderful experience. We'll have to chat sometime. Maybe some of these individuals are still around.

Dr. BRIDENBAUGH. The comment I wanted to make was really on some remark or comment that you made. It's easy to make sort of a global assessment and say we're satisfied with the outcome of these programs, but you asked a very critical question. You said: how are they being measured? And running a large R&D organization, we struggle with that every day about how to measure the effectiveness of any of the programs that we are pursuing, and you cannot do it at a kind of global level. You really have to get down and look at each specific program.

And I think one of the things that you should do, the committee should do, would be to recommend and endorse the idea that all of these programs need to have effective measures of the worth of their outcomes and make sure they are, in fact, creating value, and the Advisory Committee we've all been talking about could be a source of guidance on exactly what kind of measures ought to be put in place.

Ms. ESHOO. I think that would be a key role for the Advisory Committee, don't you? Yes.

Dr. BERLIN. I think, also, if I may just add——

Ms. ESHOO. Oh, sure.

Dr. BERLIN [continuing]. That that is a role over and above the Advisory Committee. That could be a role that is given to the National Coordination Office. The National Coordination Office has many inherent bureaucratic—there's certain things you can't do in an interagency program. People would like them to take more charge, but you just can't do that. I mean, the budgets only go—but one thing they could do, if they were given the charter, would be they could be sort of a—they could be the place that would do ongoing assessments, bring in outsiders to do ongoing assessments of the various aspects of the program in a positive sense, kind of where we are now, where we need to be going, over and above even the advisory board or in concert with the advisory board. I think it would be a very useful thing. I think they're positioned to do that, and it would be a very useful role that I think most of our colleagues would support.

I should say also, by the way, that you've mentioned, in response to Dr. Audley, that it was amazing; what you don't know is that Dr. Audley was one of those visionaries in the Government who actually was one of the early sponsors of the parallel processing program.

Ms. ESHOO. Thank you very much.

Thank you, Mr. Chairman.

Mr. BOUCHER. Thank you very much, Ms. Eshoo.

One of the reasons that we wanted to have this hearing today is because the Congressional Budget Office did release a report earlier during the year that suggested that the current focus on massively parallel computing is misplaced. Dr. Rubbert, to some extent, has endorsed that view in his testimony today, and I understand that in the aircraft industry, in particular, vector computing is more prevalent than is massively parallel computing. So I was not surprised to hear that view from him.

But I would like to get those of you who have a different view of that question to comment about it. I think a large part of the focus of the entire HPC Program for the future depends upon how we resolve this issue. And so it is an important consideration and one that I would like for other members of the panel to address in somewhat more detail this morning.

So those who have a different view from Dr. Rubbert and from the CBO report, this is your opportunity and we would welcome your thoughts and comments on the appropriateness of the focus on massively parallel computing.

Dr. Audley.

Dr. AUDLEY. Yes, sir. I'd like to observe that what we have found in our experience with parallel and distributed computing is that much of the thought process that has been directed toward massively parallel is in terms of single application, single user, single process type of host platform. What we have found in some of the applications that we use is an opportunity for many users and many processes to communicate with each other on a very random sort of basis, driven by the business flow on the trading floor, and that what we find in the large parallel system is that there are real time opportunities for computational growth and changing of the problem application mix.

What we have also found is that the computing environment seemed to be, as has been indicated, going toward ones that are more distributed with heterogeneous platforms involved in them, the distinctions being perhaps the communication band widths between the host processors and the requirement for communications between processors that are running in each of the environments.

So we find clustered workstations, single workstations on the network interacting with a mainframe that in itself is a clustered VAX machine running a database, all of these interacting with the parallel machine; that the future incorporates all of these platforms, but the massively parallel one gives the most flexibility in terms of real time expansion and growth in satisfying more users than may have been previously anticipated or scheduled without a performance degradation.

Mr. BOUCHER. So you, then, would welcome the current focus on massively parallel computing in terms of its application in the securities industry?

Dr. AUDLEY. It has been an appropriate level.

Mr. BOUCHER. Okay, good.

Dr. Berlin, would you care to comment? I think you have addressed this issue to some extent. Do you have anything to add?

Dr. BERLIN. Yes. I think there's two things that I'd like to say. One is that the issue is no longer—for the next—as we look forward, we can argue about whether decisions were made properly looking backward, and, frankly, there were risks taken, and we learned some valuable lessons. I don't think we've admitted them sometimes, what we've learned, but we've learned that these hard things to do. They're very hard problems and we've learned some of the difficulties involved, and we've learned that it's going to be a long time before massively parallel processors are just a general purpose computer. And so we've kind of—it kind of worked out for us.

But I think as we look to the future, we need to declare that the issue of whether it's massively parallel or not massively parallel is not the issue. The issue is what the user wants and what the user needs.

And I would like, for one—and I know I have a lot of people who would support this position—I would like to see the program start to give more money directly to the users and take the model of the Japanese robotics industry. When the Japanese set robotics as a major focus, they didn't settle on a few architectures and pick them and then go with them. What they did is they got things started and then they started giving money to users, and users went and they picked the ones that were meeting their needs the best. And when those—when the ones that were left standing came to the United States shores, we said, my gosh, what strong competitors; that must be Japan, Inc. It wasn't Japan, Inc. at all; they just fought to the death already and they came over here having met the needs of the users.

So as we look forward, I think really we need to focus totally on what the users want because there's no market for parallel processing, but there's no market for vector processing. There's only a market for a job to be done.

Mr. BOUCHER. Thank you very much.

Dr. Bridenbaugh.

Dr. BRIDENBAUGH. Well, it's a little hard to add anything to what Dr. Berlin and Dr. Audley have said. I agree with them.

I think a bit of your view depends on the nature of the kind of problems that you are trying to solve, and the problems that we face seem to be more suited to parallel processors than not. We believe that distributed computing in all its forms is the wave of the future. It's maybe even the wave of right now. So I would—I think the focus is about appropriate. I would agree that it probably needs to have more pull from the customers or from the user base, as Dr. Berlin said.

Thank you.

Mr. BOUCHER. One of the things that the report indicated is—and I'll just quote this—it says that, in referring to parallel computers, "The class of computers has yet to demonstrate their utility as general purpose machines. They are difficult to program and use and they have too few applications."

Yes, Dr. Bridenbaugh, your response?

Dr. BRIDENBAUGH. Well, I think that just points out the problem, and in my written testimony and what I tried to say earlier is that the real barrier is availability of software right now, and that's right where we should be focusing our effort, to produce more application software for parallel machines and then you'll start to see the outcomes from the earlier expenditures.

Mr. BOUCHER. Can we contest successfully the notion that they are not general purpose machines? I mean, we see some pretty impressive applications just represented on this panel today from the petroleum industry, from the securities industry, really divergent sectors of our economy. Are they becoming general purpose machines, at least as reflected by those suggestions?

Dr. Audley.

Dr. AUDLEY. I think they have demonstrated a utility. I would say that they are not general purpose machines at this date. There's more work to be done. I think they show promise and I think that the follow-through should be pursued to see if they can fulfill that promise.

As far as programming the machines, in our experience, as I mentioned earlier, we had three criteria. One was performance; the other was that we didn't want to make our applications programmers that we had in-house obsolete by bringing in some new technology in 1988 when we brought in the first machine. Actually, what we found, though, was quite the opposite; they were delighted to have new high technology to take a look at. Most of them took very well to it. Over time, all of our applications programmers developed applications for the parallel machine as part of the suite of platforms that we use.

Mr. BOUCHER. Dr. Rubbert, I'm going to come to you last because I'll give you rebuttal time, but let me hear first from Dr. Frazer.

Dr. FRAZER. Well, I think from my testimony that my bias in this area is well known. I believe that massively parallel machines have much broader applicability than vector machines. I've built vector machines in my career and I'm somewhat familiar with their capabilities.

I think in looking at the focus of the HPCC Program overall, the one minor criticism that I would have, which nonetheless I think is urgent, is that there's been a tendency to focus on problems and areas in which the computing problem is described as one of arithmetic. In our particular industry—that's why I tried to focus on nonarithmetic processing, nonnumerical processing—the database search problem, the text problem, the video server problem, the E-Mail problem, the credit card and debit card processing problem, none of those are Grand Challenge mathematically formidable problems, and yet they are major problems confronting our industry across the board.

I believe that without question massively parallel technology has a much broader applicability to that class of problems which has historically not been the focus of HPCC than has yet to be appreciated or realized.

Mr. BOUCHER. Okay, thank you.

Dr. Bloomquist.

Dr. BLOOMQUIST. Well, first, let me just say I think people are a little misguided when they want to think of these as general purpose computers. A general purpose computer today is sitting on people's desktop. It no longer sits in the computer center. A vector, large vector computer is not a general purpose computer anymore either. It certainly has a limited range of applications. You don't see it being used in database, for example, whereas a massively parallel computer can.

Certainly, a massively parallel computer does have a broad range of applications that will be developed in the future. We're fortunate in the petroleum industry to have some applications that run very well on these machines today. We can get performances that are 50 times that of our Cray on certain applications, certainly not all applications, but you only want some very large applications to run on these types of machines anyway, and we certainly feel like this is the leading edge; this is the area where you're going to get an order of magnitude improvement and competing in a short period of time. You're not going to see those kinds of improvements with conventional vector supercomputers, and you've potentially even got the potential of an order of magnitude, two to three orders of magnitude, and these open up a whole new set of problems you can solve. It's not just doing the things you do today faster. It actually opens up a new horizon of applications that you hadn't even been able to think about in the past.

Mr. BOUCHER. That's a very interesting point. And so what you're suggesting is that the statement in the CBO report that this class of computers are not general purpose machines is really an irrelevant criticism within the context of whether they deserve this much focus in terms of the HPCC Program. They have special applications that suggest great utility for the society, nonetheless.

Dr. BLOOMQUIST. This whole idea of the whole world's a network, this is one of the nodes on the network that does very large problems and that's what you want to use it for. Your general purpose machine is really one that sits on your desk, the workstation or the personal computer that you use to do a wide range of problems, including analyze data that comes off of these big machines.

Mr. BOUCHER. Okay, Dr. Rubbert.

Dr. RUBBERT. Yes. Well, our perception today is that massively parallel machines are niche machines that have in certain cases very, very good application. We have one on low observable calculations, radar waves bouncing off of vehicles. It does that better and cheaper than anything else. That's a niche. But our perception is they're not general purpose and they are hard to program.

Now in support of our commercial transport aircraft development, we are very active in continuing to develop software and we are closely and strategically linked to the long-range business plan of the company. How we make our plans for what we're going to do and not do in computational development is we look ahead at our business plan and in a certain year we hope to launch a new airplane program. We back up from that and say here's the year we're going to be designing it; what are the principal technical and design issues that we can improve upon? And then the time between now and then is to do the development to do that. So we're moving into the world of what I would call just-in-time application software development keyed directly to long-range business plans of the company.

We can't predict what the particular issues will be. One airplane, it may be this issue; another airplane, it may be something else. So we also need general purpose. We need general purpose capability and just-in-time—and the ability to do easy software development so we can play the just-in-time game. Today the massively parallel machines do not seem to do a very good job of supporting those two requirements.

Mr. BOUCHER. You would not suggest that we disinvest in massively parallel computing research, however. It sounds to me like what you're suggesting is that we ought to also focus on other kinds of——

Dr. RUBBERT. Yes.

Mr. BOUCHER [continuing]. Computing applications.

Dr. RUBBERT. We need balance in our program.

Mr. BOUCHER. Is that a fair summary of your suggestion?

Dr. RUBBERT. Yes.

Mr. BOUCHER. Okay, very good. Thank you. I do ask good leading questions. Those were the objections that my opposing counsel always succeeded—I was well known at the local bar for asking leading questions. [Laughter.]

Let me just ask one other question of this panel and then we'll conclude this morning's hearing. Well, I guess Mr. Minge's gone.

The Administration asked for a billion dollars in funding for the HPCC Program for Fiscal Year 1994. The Appropriations Committees have produced bills in the House and Senate, still as of this point unresolved, that at the worst case would produce about a 13 percent increase. The billion dollars is about a 39 percent increase over Fiscal Year 1993.

Taking the worst case that might come out of the appropriations process, we'll get about a 13 percent increase over Fiscal Year 1993. I am personally somewhat concerned that if the worst case comes about this program will not advance as had been anticipated. And I wonder if this panel could give us some sense of what the general risk is in terms of a slowdown of the pace of development of the HPCC Program and of the pace of funding. Is there some risk to our international competitive position, not just in the field of the technologies that we're performing research for, but in the various fields of the users of these technologies? What kinds of risks do we run if we slow down the pace of funding for the HPCC Program? Any volunteers for that?

Dr. Berlin.

Dr. BERLIN. We have discovered in the last 10 years that we're facing some very hard problems to get to the next level of broad application. We have also discovered that the rewards are unbelievable.

What Boeing has done, every time you take off in a Boeing 767 or 757—I believe those two were the ones—it saves fuel, every time it takes off for the entire life of the airplane. That's just one of hundreds of examples.

Now the problem is that when we started this journey—for example, the auto industry was talking about, could they shave one year out of three off of the design turnaround, but now we're looking at design turnarounds in some products that are a matter of months, where whole market shares change.

So there's two risks that I see. One is we will—we need to get there sooner rather than later across the board, hybrid computing, parallel computing, vector computing. Secondly, the largest benefit of the program has been the spinoff along the way. And the thing that has made this program unique is it's gone on this growth curve that it has enjoyed, is that the spinoff for industry such as Alcoa—and for every Alcoa there's 100 more stories like that—the spinoff as we go along the way has been a churning of ideas and thoughts and innovations that has been critical of the entrepreneurial process as we go through this major change of thought process in the whole way the world does things and we go to optimized products as opposed to functional products. That's a very major change in the whole thought process, and as we train the pipeline of people to think that way, we are fueling the engine. And by slowing that down, I believe this is one area that there's a great risk. This is not just a niche technology. This is where it's at. It's what it's all about.

Mr. BOUCHER. All right. Dr. Rubbert.

Dr. RUBBERT. Yes. I would like to support, reinforce what Dr. Berlin had to say. One of his early comments was the influence on cycle time. It's one of the two stated corporate goals of the Boeing company at the present time, is to within the next three years, starting from a certain date in 1993, to cut the cycle time between order and delivery of an airplane in half. That's what we're aiming to do. The other goal is to cut the cost of that airplane by 25 percent within the next five years. And I feel fairly confident that we will meet those goals.

We've got to learn to look at the power of computing not in terms of what can you do, what wonderful problem can you do today that you couldn't do last year, in the direction of how much faster can you do it and that's where we really see payoff. You know, wind tunnel-based design process is horrible. It takes months to build a wind tunnel model, months to schedule the tunnel, et cetera. Computing something is a matter of hours rather than months. We have to learn to look at computing as a speed machine rather than just a sledge hammer to solve problems we couldn't do before, and that's where you're going to find a great improvement in industrial competitiveness in this country.

Mr. BOUCHER. Well, those are excellent responses to that question. Let me go from that more specific question finally to a more general one, and this lies basically at the root of all the debate that we've had on the high performance computing program.

When we had H.R. 1757 on the floor of the House, this debate was reflected once again, and I'm confident that it will be as long as the Federal Government is investing in new networking technology and high performance computing. The argument against this effort goes something like this: that the private sector, left to its own devices, will do these things anyway; that the Federal Government really does not have a role in providing research and development funding for new networking technologies and doesn't have a role in providing research and development funding for a new generation of high performance computing, and that the Government's best role would be to simply stand aside and let the private sector do all of these things entirely on its own.

Are the proponents of that argument correct? Is this a Federal effort that best would be left aside or are we pursuing the proper role? And if we are pursuing the proper role, tell us why this is something that would not be done without the Federal Government's aid and assistance through research. Dr. Rubbert?

Dr. RUBBERT. Yes. I am a firm believer in market forces leading to the best type of products that are most suited for us. I view the role of the Government as working in a way that enhances and supports the working of those market processes rather than replacing a market-based decision on some scientist's decision on what the world needs. The idea of getting the customer strongly involved upfront in the initial planning, et cetera, and working hard to get the computing industry talking to the applications industry, so we get the earliest view of what's coming down the pike in computing and they understand our requirements better—in other words, use government involvement to enhance the working of the market-based process, and I think you win that way.

Mr. BOUCHER. All right. Well, I agree with that test, and I think we all would. But using that test, how do you apply it to what we're doing in terms of high performance computing and high speed networking? Are we on the right track or are these technologies that the private sector would develop in the absence of a Government role?

Dr. Rubbert, would you care to respond?

Dr. RUBBERT. Well, that's a broad question that goes beyond my area of expertise, but things like the national network, and so forth, it seems obvious to me that that's a clear federal role, just like the highway system that we drive our cars on.

Mr. BOUCHER. Okay, thank you. Dr. Bridenbaugh, did you want to respond?

Dr. BRIDENBAUGH. Well, I would agree with his comment that the—certainly the network is an appropriate role for the government. I am not sure that I—my understanding of what you said earlier was you're not talking about slowing down the program; you're talking about slowing down the growth of the program; right? And I think that in hard times, which almost all industry in the United States is in right now, we find a way to do things with less technically and scientifically. And so I would be a little on the side of I'm not so sure that the difference in those two numbers is significant enough to cause us great amount of concern, and the fact is I think maybe people find more creative ways to get the work done.

I don't know how you would get all the American industry to step up to some of these big programs in some kind of unified, "we're all going to act together" approach to addressing these issues. So I think it's a role of the Federal Government to do that, and I think it's appropriate in this whole—the role that you're playing in this whole area of supercomputing is appropriate.

Mr. BOUCHER. Okay. Further comments? Dr. Frazer.

Dr. FRAZER. I think it seems to me there are two areas in which the Government can legitimately play a role. One is in the stimulation of, what I call, seed technologies, which MPP was when the Federal Government began to invest in it. And the second is in investing in areas of infrastructure, such as the information highway, in which it's, No. 1, totally infeasible for industry to undertake the kind of collective action that would be required in order—under present law, as I understand it, the kind of collective action that would be required without some overall coordination role.

And, secondly, without some kind of stimulation pushing that forward, you fly in the face of what has historically been the investment and amortization cycle of the telecommunication industry, which is 30 or 40 years to depreciate equipment, and, therefore, there's an entrenched resistance to having to go to new technology rapidly.

So I do believe that the government can play a role in getting technologies to the point where the natural competitive forces can take them forward. In the area of high performance computing particularly, that requires a large investment initially. It isn't something like the personal computer where, with a year or two of investment, a broad market opens up and one can see a return. It's taken a number of years for MPP to get to the stage where it is now, and it's entirely appropriate for government to focus on arenas such as that, it seems to me.

Mr. BOUCHER. All right. Dr. Audley.

Dr. AUDLEY. I would observe that Government has a responsibility for the leadership in many arenas, certainly to participate with the population, organize and manage the society. I think what we normally might expect from industry and commerce is to deliver the fruits of that organizational and leadership process. I don't think that we should totally decentralize the leadership role.

Mr. BOUCHER. All right. I want to express our appreciation to the subcommittee—of the subcommittee to this panel of witnesses for their attendance here this morning. It's very good to hear from you.

I can assure you that we will encourage the Administration very expeditiously to appoint the external Advisory Committee, which everyone has indicated a great need for, and, hopefully, that will happen very shortly.

We may be calling on you for additional advice. In fact, we probably will as we continue to oversee the development of the High Performance Computing Program.

Your testimony has been very helpful to us and we'll look forward to your future guidance in other avenues and hearings. Thank you very much.

And the hearing is adjourned.

[Whereupon, at 11:44 a.m., the subcommittee adjourned, subject to the call of the Chair.]

○

75-413 (11?)

ISBN 0-16-043524-2

90000

9 780160 435249