

# CISE NSF Centers for ASC division

Catalog number: 102740245  
Gordon Bell NSF-CISE Records  
Virtual Bankers Box<sup>1</sup> 3



---

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

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



## Advanced Scientific Computing: Past... Future

### **Phase I.** Supercomputer Access for Research Community (9/84-10/86)

*Supply supercomputer power to existing scientific and engineering users.*

Capacity: Universities [3(4);1(1);0]\*; NSF [0;2(2);0] = 5; 2

Access: basic terminals via dial-up lines

### **Phase II.** Establish Supercomputer Centers (1985-

*Establish centers to provide computational power and begin to train new paradigm for computation. Establish a network to link centers and provide users high speed access with graphics workstations.*

Capacity: Universities [5(7);1.75(14);40(40~80);1(4)]; NSF [2(3);5(33);1(4)] = 69; 40

Access: terminals, Local Area Networks, and workstations

Network: NSF backbone @ 56 Kbps; Consortia nets; dialup

Training: Continuous on campus training, Summer institutes for all disciplines

Environment: graphical i/o via workstations; start to hide-the-supercomputer for improved ease-of-use; begin to provide a "standard" environment for operating system, languages and graphics across workstations, mini-supercomputers, and supercomputers

### **Future** Distributed, Compatible, and Visual Supercomputing ( $\geq 1990$ )

*Provide the leading edge environment with "visualization" where users can compute at any machine in a fully compatible hierarchy depending on cost, performance, and geographic needs. Initiate a Computational Science and Engineering program. Distributed communities of researchers working on common problems.*

Capacity: Universities [4(80); 8(160); 100(200)]; NSF[1(20);4(128);1(12)] = 480; 160

Access and Network: Research Network (Global LAN) based on fiber optics packet switching operating at 45~140 Mbps to interconnect entire research community.

**Peak Power:** By using parallelism provide X10 ~ X100 speedups

*By providing order of magnitude power increase, new problem solutions will become tractable.*

Notes:

\* [number of supercomputers by CDC; Cray; IBM; Japanese (Amdahl/Fujitsu; Hitachi/National; NEC)]. () power in equivalent Cray 1's.

## **Advanced Scientific Computing Program Evolution**

### **Original OASC Program Goals**

- supply supercomputer power to scientific and engineering community
- train scientific and engineering community to use supercomputers
- support U. S. supercomputer industry
- provide for remote access by users
- stimulate development of a "rich and powerful" scientific/engineering computing environment **for all** as measured by total power and utility*

### **Evolving Centers Program**

- provide the highest performance service with the latest supercomputers, acting as "beta" sites for the manufacturers
- provide generic applications software libraries
- maintain "consortia nets", and connect to NSFnet and regional nets for user access
- train "selected or strategic" users in various disciplines (30/center)
- allocate time; report use by disciplines, user experience, user job size, and geography; review effectiveness for users; and review final scientific/engineering output
- co-ordinate and establish a common program, user interface, and graphics environment across centers to maximize user effectiveness ?*
- establish key links with hardware and software suppliers?*
- establish a program for industrial users?*

### **CISE-wide Programs in CSE, New Technologies, and DCCR**

- benchmark and understand performance and cost-effectiveness of various conventional scientific, multi-vector processor computers
- enhance training for all disciplines to include vector processing
- provide advanced graphics techniques (visualization)
- provide parallel processing environments and encourage applications appropriate to efficient use of multiprocessor supercomputers
- research on advanced algorithm and software across disciplines
- research new technologies for potential performance breakthroughs

### **Other NSF CSE Programs**

- develop and maintain area(x)-specific, common user communities for common programs and datasets
- sponsor and organize the "grand challenges" in computational science

### **Changing Networking Activities-NCRI**

- Originally provided network to link centers to users, to each other, and to other networks.
- Currently links to other databases, going beyond traditional computing
- Extend network to support entire research infrastructure*

Gordon Bell; 10 March 1987

## **Advanced Scientific Computing Program Evolution**

### **Original OASC Program Goals**

- supply supercomputer power to scientific and engineering community
- train scientific and engineering community to use supercomputers
- support U. S. supercomputer industry
- provide for remote access by users
- stimulate development of a "rich and powerful" scientific/engineering computing environment for all as measured by total power and utility*

### **Evolving Centers Program**

- provide the highest performance service with the latest supercomputers, acting as "beta" sites for the manufacturers
- provide generic applications software libraries
- maintain "consortia nets", and connect to NSFnet and regional nets for user access
- train "selected or strategic" users in various disciplines (30/center)
- allocate time; report use by disciplines, user experience, user job size, and geography; review effectiveness for users; and review final scientific/engineering output
- co-ordinate and establish a common program, user interface, and graphics environment across centers to maximize user effectiveness*
- establish key links with hardware and software suppliers?*
- establish a program for industrial users?*

### **CISE-wide Programs in CSE, New Technologies, and DCCR**

- benchmark and understand performance and cost-effectiveness of various conventional scientific, multi-vector processor computers
- enhance training for all disciplines to include vector processing
- provide advanced graphics techniques (visualization)
- provide parallel processing environments and encourage applications appropriate to efficient use of multiprocessor supercomputers
- research on advanced algorithm and software across disciplines
- research new technologies for potential performance breakthroughs

### **Other NSF CSE Programs**

- develop and maintain area(x)-specific, common user communities for common programs and datasets
- sponsor and organize the "grand challenges" in computational science

## Advanced Scientific Computing Needs and Opportunities

	FY 88	89	90	91	92
Centers Today					
<b>Current Plan</b>	<b>45</b>				
<b>Co-operative Agreement</b>	<b>48.5</b>	<b>51.9</b>	<b>54.6</b>		
Minimum staff upgrade	1.5	1.7	1.7		
Cornell second 3090/600	2.5	2.5	2.5		
Misc upgrades at 5 yr. replace	1.6	6.9	6.9		
Total if 5 year replacement	53.7	63.2	63.9		
Total if 3 year replacement	59	73	78.9		
Centers with more capabilities	54	63	69	74	80
Enhance range of activities	2	2	2	3	3
Enhance equipment at centers		10	29	33	37
Total Centers Plan	56	75	100	110	120

### **CSE & Technologies Today**                   **3.6**

CSE, Technologies, Visualization, etc. for a "balanced" program

CSE	6.5	8.6	10.7
New Technologies	3.6	4.6	5.4
Visualization (center/group)	4	5	6
Sci and Tech. Centers	3	5	6
Total	17.1	23.2	28.1

Support purchase of Mini-supers

for groups on matching basis	0.25/group, one time equipment grant		
machines	40	48	56
cost	10	12	14

# **Preparing for Changing Scientific Computing Environments**

**Gordon Bell, Assistant Director  
Computer and Information Science and Engineering  
National Science Foundation  
Washington, D.C. 20550  
30 September 1986**

## **Introduction**

Recently, a hierarchy of scientific computers in three price ranges and computing styles have evolved with relatively the same performance/price and computational ability. The hierarchy includes: the supercomputer and large mainframe used as a regional or central computer costing between \$10M-20M; the mini-supercomputer used alternatively as a central, departmental, or group computer costing around \$500K; and a workstation/workstation cluster, used as a shared, departmental resource, as a single user system, and access to other machines in the hierarchy costing around \$50K.

The comparable computational power of these new scientific computers raises various policy issues for NSF including the management of its Advanced Scientific Computing Program, the role of the five National Centers, and the way computation is supplied to the research community. Ideally, a user will utilize all forms of computation based on economics, networking, power, response time, and interaction (especially graphics) needs. This paper explores these parameters and outlines the policy implications required to provide the most productive environment for the research community.

The data for this analysis are key performance characteristics of a variety of scientific computers:

- number of processors, #P.c
- primary memory size in 64-bit Megawords, M.p, with virtual memory (shown as .v)
- secondary memory size in Megabytes, M.s
- speed measured in millions of floating point operations per second using Dongarra's Linpack benchmark for a 100x100 and 300x300 matrices, Mflp
- the price of the machine in millions of dollars, \$.M
- the cost-effectiveness, i.e. performance per unit price for two sized matrices, flp./\$
- introduction date, Intr
- stretch time versus Cray XMP single processor for a single job ()

**Table of Computer Characteristics**

System	#P.c	M.p	M.s	Mflp	\$.M	flp./\$	Intr	(Stretch) comments
<u>Supercomputers</u>								
Cray 416	4	16	9.6	108.-480	17	6.4-28.2	86	(1) 27/Pc, 8.5ns clock
Cray 48	4	8	9.6	108.-480	15	7.2-32	84	
ETA 10 (Est.)	8	288.v	9.6	1040.-2K	19.7	52.8-107	6/87	(0.2) 10.5ns, 7ns '88 = x1.5
Cray 2	4	256	9.6	60.-372	18.6	3.2-20	86	
<u>Mainframe</u>								
IBM 3090/400	4	16.v	60	48.-108	9.8	4.9-11	9/86	(2.25) sans software
<u>Mini-supercomputers</u>								
Alliant F8	8	1.v	.4	7.6-14	.75	10.1-18.7	6/86	(3.6) with directives
Convex C1	1	1.v	.4	2.9-14	.4	7.3-35	1/85	(9.3)
SCS-40	1	2	.7	7.3-26	.65	11.2-40	7/86	(3.7) XMP compatible
<u>Workstations</u>								
Sun 3-200	1	1.v	-	.47	.04	12.0	9/86	(57)
Sun 3-200	1	1.v	.28	.47	.06	8.2		(57)
Sun 3-200	1	2.v	2	.47	.12	3.9	9/86	(57)
+3 diskless	4	8.v	2	1.9	.25	7.5		cluster of 4
Sun 87/B Joy	1	2.v	-	1.5	.02	75	87	(18)
Sun 88/B Joy	1	4.v	-	4.0	.03	132	88	(6.75)
Sun 89/B Joy	1	8.v	-	10.0	.04	250	89	(2.7)
<u>Historical References</u>								
Cray 1/S	1	1		12.-66	6	2.-11	75	(2.3)
VAX-11/780	1	.5v	.1	.15	.3	2.	4/78	(180)

Notes:


For the Crays and ETA-10, the performance is for 4 and 8 independent job streams. Linpack appears to be a good benchmark in that it correlates well with other scientific and engineering benchmarks, and with the average delivered power. As work is "tuned" for vector processing, the 300x300 matrix is a realistic target for typical applications.

A variety of different secondary memory configurations are given, including none for the 3090 and several Suns.

The fastest uniprocessor is the NEC SX-2 at 43-347 Mflops.

Super-minicomputers, and high-performance PC's are not included because they provide relatively poorer performance, and performance/price. For example, the PC AT/370 is a factor of 818 slower than a Cray and the cost to perform floating point operations is roughly double.

# Comparative Linpack Performance



Source: Argonne National Laboratory  
 Technical Memorandum No. 23  
 September 7, 1986  
 Jack J. Dongarra

## **Background**

The current surge of interest in supercomputers becomes clear when we look at the evolution from the late 70's when the Cray 1 and VAX 780 were the standards for computation. The 780 entered the scientific and engineering community because it provided relatively the same price performance as a Cray 1, even though the performance differed by a factor of 80 (using Linpack as an indicator). A more reasonable estimate for the difference is more like a factor of 20-40. Those who bought VAXen observed that since the average user only got 1-2 hours of Cray time each week, (50-100 hours per year) they could get the same amount of computing done by letting a VAX grind 20-160 hours per week.

Over time, the Cray evolved; the XMP was speeded up by over a factor of two and built as a multiprocessor, which roughly trebled the performance/price. When the scientific community started utilizing Crays with improved compilers, they began to develop more effective algorithms for vectors that increased the effective power of the machines. The delay in getting a more cost-effective VAX (the 8600 was two years late), and the relatively high price of VAXen exacerbated the difference between the supercomputer, and the super-minicomputer (in essence a lower priced mainframe). The popularity of VAXen for more general computing also allowed the price to remain high, by giving it a market outside the research community. DEC, like IBM when it introduced a complete range of compatible computers, may have become less interested in and attentive to the research community. The Cray/VAX gap may have been a major motivation in the formation of the NSF Advanced Scientific Computing Program.

In the early 80's Alliant, Convex, and Scientific Computer Systems formed to exploit the performance/price gap between the Cray XMP and VAX by utilizing vector data-types pioneered in the Cray 1. Thus, a new class of mini-supercomputers was formed, all of which have better performance/price than the Cray (almost a factor of 2 in the case of the new SCS-40).

By 1985, ten years after the Cray 1, IBM and Japanese manufacturers building IBM-compatible mainframes had added vectors and multi-processors to their machines.

## **Observations About the Computers From the Table**

Three characteristics are important: the processing power in Megaflops; the cost-effectiveness in flops/\$, and the stretch time versus a Cray. There are exceptional computers, when comparing the cost-effectiveness in each class: the (projected) ETA-10 (to be better by a factor of 8!), and the SCS-40 (better by almost a factor of 2). The SCS-40's virtue and principle flaw is Cray compatibility. Other mini-supers have virtual memory. A cluster of SUN workstations could provide up to a factor of 2 better performance/price, depending on the amount of secondary memory. The factor of 5 difference in the speed of the ETA-10 versus a Cray XMP should open up new problem solution domains. The ETA-10 uses large CMOS gate arrays on large, multilayer



printed circuit boards. This kind of fabrication provides a potential breakthrough in cost that is counter to the use of ECL to build supercomputers, large mainframes, and superminicomputers by Cray, DEC, IBM and the Japanese. Both the Cray 2 and ETA-10 have large memories that should open up new problem domains. All of the machines, except the Crays, have virtual memory. Because of the lack of paging, it may be difficult for multiple users with very large problems to effectively utilize the Cray 2. The use of large physical and virtual memories needs to be explored and understood.

While the table shows times for a floating-point intense program, Linpack, it is unclear how the machines perform under comparable workloads or whether they will actually be used in the same fashion. For example, a slower machine is likely to be used more interactively and results of the computation viewed constantly to avoid unnecessary work. Users of large batch machines may have to request more work and output because turn-around is longer. Scalar benchmarks aren't given, and most machines are used a significant amount of time either interactively or in scalar mode, both of which lower the performance and favor the 3090 (which outperforms the Crays in scalar mode), mini-supers, and workstations.

NEC's SX-2, not included in the Table, executes Linpack at about twice the performance of a single processor Cray XMP. The performance/price is unclear.

Many computers exhibit performance/price comparable to today's supercomputers. The Advanced Scientific Computing Program must understand the relative power and work capacity of all forms of computation and begin to develop ways to supply resources appropriate to user need and cost-effectiveness considerations.

### **Can Users Tolerate the Time Stretch/ Lower Cost Trade-off?**

Can a user of a smaller computer, stand the lengthened turn-around time that comes with using a slower computer and stretching the computation time by factors of 4 to 10? At present, only one or two users within our user community are receiving an hour of computer time per day. The mini-supercomputers, supplying the equivalent of one hour of Cray time in 4-10 hours are competitive because the average turn-around for a one-hour job on a Cray can easily be this long. The typical turn-around for a 15 minute job is 2 hours (or factor of 8 stretch). The Sun Workstation might be used for longer computation provided the user "guides" the computation. The Sun's stretch factor is comparable to that experienced between the Cray and 780 during the late 70's. Alternatively, advances in partitioning programs for parallel processing make the cluster have the best performance/price if a job can be parallelized using a message-passing model of computation.

Based on the performance, and time allocations inherent in supercomputer use, a complete

hierarchy of computers will exist and is justified. Given that an individual user or project is likely to simultaneously access all levels of the hierarchy, a compatible (and most likely standardized) basic environment that can support user communities, who in turn have common applications environments, is essential.

### **Multiprocessors, Array processors and Multicomputers (e.g. Hypercubes) for Parallel Processing**

A number of alternatives exist that may offer significant improvements in performance or performance/price. For example, a 64 computer NCUBE has been used to solve a problem that took twice as long on a single processor XMP. The improvement yielded almost an order of magnitude in cost. Given the decomposition for parallel processing on the NCUBE, an XMP might be used to gain a 4 times speed-up; in fact, the XMP operating in this mode has computed Linpack at a rate of 713 Mflops which is 26 times the single processor rate. Likewise, array processors such as the FPS X64 have been lashed to minis and mainframes, yielding significant improvements in performance/price. None of these alternatives are explored.

Standardized parallel processing primitives in all programming languages based on a multi-process, message passing model of computation is needed for all structures. Programs used in this fashion will operate compatibly and identically across workstation clusters, multicomputers such as the hypercube, and shared-memory multiprocessors (e.g. Cray and ETA). Given the relatively constant performance/price and similar turn-around times for all of the computing alternatives, parallel processing becomes essential.

### **The Role of the Super Computer Centers**

Historically, centers have existed for a variety of reasons including cost sharing, technology, performance, networking, user needs, local politics, government funding, etc. Clearly when hot ideas emerge and projects need ten to several hundred hours of supercomputer time that can't be supplied locally the centers are essential. The definition of the kinds of work that the centers will support is critical, given that computation can be done very effectively by local university centers, departments, projects, and individuals at workstations.

Our centers are critical to scientific and engineering computing for the research community. Today the centers train users about the parallelism inherent with vector data-types. They have the programs and staff to train the trainers and users rapidly, and to support large programs and datasets inherent in supercomputer use. Centers may be the best place to support certain large programs and databases for a given intellectual community; NCAR is an excellent example as it provides millions of lines of common programs and 17 terabits of common data for its community of atmospheric scientists to environmental engineers. Centers may also support common programs for communities of distributed users at mini-supers, super-minis, and workstations in order to

supply service when the distributed research requires significant computing power.

Large amounts of power (on the order of 1 hour per day) would be supplied to large projects that do not have machines, and to a community of student and casual users who access common programs and data. If the "average" project uses 1 hour per day or 350 hours per year, then a Cray XMP would support 24x4, or about 100 projects! Projects of this size would be, in effect, subsidized at about \$100,000 with steady-state costs. It can, alternatively, service 640 users who use at most an hour a week, or 50 hours per year, providing them about a \$15,000 subsidy. Finally, several thousand student and casual users who would use no more than 10 hours per year (a year on a PC/370) could be supported at negligible cost. Policy statements are needed which characterize useage across geography, user size, and discipline.

The centers have a lead role in supporting state-of-the-art computers of all types including supercomputers, mini-supercomputers, and larger scale experimental machines. The centers should be the beta test sites of all new systems, especially those which can not be easily purchased or supported by local researchers or departments. The centers must take the lead role in understanding benchmarks, workloads, and cost-effectiveness of all forms of computation.

**Standards.** The three alternative forms of computation that form the main line of computing all provide roughly the same computational service at comparable costs (not including the cost to the user). We must establish standards that make it equally easy for users to work at any of the places in a compatible fashion. In many cases, a user will use the super or mini-super or existing super-mini for calculations and the workstation to view results. Thus code will be run in a highly distributed fashion across different machines including new, and evolving UNIX-compatible PC's. Similarly, we should work toward establishing and supporting common programs and data across engineering and scientific disciplines so they may compute at any level of the hierarchy.

## Conclusions

Computers now exist which allow various styles of computing ranging from regional supercomputers to personal workstations. All of the computers in the hierarchy will continue to exist and flourish because, with the exception of the ETA 10 to be delivered next year, all offer relatively the same cost and effectiveness.

Having the wide range of styles and locations demands attention to:

- training, education and program support;
- networks for intercommunication of programs, data, and terminal access;
- benchmarks, workloads, accounting, and pricing i.e. understanding cost and effectiveness;
- allocation of time across user communities by size, discipline, and geography;
- standardized programming environments and graphics enabling effective use;
- supporting specialized community programs (e.g. NASTRAN) and databases (e.g. NCAR);
- specialized and alternative computers; and
- standards, understanding and training for compatible, message-passing parallel processing.

With the center program entering phase II, attention and resources will have to be focused on these demands.


# NOMOGRAPH

## of Computer Simulation

**COMPUTER  
RESPONSE TIME  
(SEC/PROBLEM)**

**SUSTAINED  
EXECUTION  
RATE (OP/SEC)**

**PROBLEM  
COMPLEXITY  
(OP/PROBLEM)**



## MEMORANDUM

SUBJECT: Advanced Scientific Computing Strategic Opportunities  
TO: FILE  
FROM: Assistant Director/CISE  
DATE: August 31, 1987

I'd like to raise several strategic issues for our support of advanced computing for research which go beyond the budgeting for the NSF Supercomputer Centers. This concern is based on:

- o lack of support for distributed minisupercomputer facilities which complement, but do not replace the DASC supercomputer centers
- o difficulties and long lead times in the creation and support of a completely compatible distributed computing environment across a range of machines horizontally (supers) and vertically (supers, mainframes, super-minis, mini-supers, and workstations)
- o inability to support new, faster and/or more cost-effective computers on an opportunistic basis, (because the centers consume such a large fraction of our resources)
- o the inefficiency of central facility versus research discipline-based management
- o inability to support a CSE or a "Grand Challenges" program via the research directorates (which is necessary for a true revolution in scientific computing)
- o lack of understanding and inability to fund benchmarks dealing with performance and performance/cost issues associated with alternative forms of computing

### Background

Last fall, I wrote a memo which was circulated widely within the community (Preparing for Changing Scientific Computing Environment, 30 September 1986). The reactions were bi-modal: "don't rock the boat because at least we have supercomputer time" and "you're right, we need a range of compatible computers for a complete environment, furthermore we need support for

smaller machines". I am now convinced the hypothesis of the Cray XMP/VAX-780 gap was correct, and the genesis of the supercomputer program. However, if the program hadn't been started, then today, more researchers would be using mini-supers and super-minicomputers such as Alliant, Convex, DEC, Elexsi, IBM, etc., which would serve many researchers better than the centers are able to. For those researchers who do not need the maximum capability offered by supercomputers, mini supercomputers distributed into the researchers' own environment and connected via networks to the supercomputer centers would offer significantly more and higher quality computing resources to those researchers. Note that capacity refers to throughput and capability refers to turnaround. Independently, Stuart Rice made similar points in a recent speech:

"I have in mind a networked system . . . graphics workstations and local supporting intermediate computer and ultimately connects to a supercomputer, with provisions of special devices . . .

1. Distribution of computer resources distorted by the use of "funny money" . . . cash and credit . . . Workstations come from grants, supers are "free" . . . intermediate machines are indispensable and current funding patterns have to change if . . .
2. Dramatic advances in hardware haven't been matched by advances in algorithms and operating systems. . . . parallelism is "chicken and egg"
3. The scientific community has become rather inflexible with respect to use of operating systems. . . . don't use particular machine features . . .
4. . . . the scientific community has not been as imaginative as it might in thinking about the uses of computation in research."

### A Vertically and Horizontally Compatible Computing Environment

Rice's scenario is based on complete horizontal (all the supercomputers) and vertical (supers, mini-supers, and workstation) compatibility. We now have six incompatible centers (including NCAR). Since UNIX is truly becoming an important standard under the forthcoming Federal Information Processing Standard, FIPS P1009, Posix, I believe this situation must change. Unfortunately, NSF at SDSC is supporting the development of "Unix-like" calls for CTSS, thereby creating a continuing support commitment for a proprietary operating system. The ASC Program has been slowly addressing this problem. Until UNIX is in place the existing situation limits the use of the centers' supercomputers for entire communities such as VLSI designers and creating more work for their users who now operate with the standard (workstations and minis).

## Mini-supers Are Required to Operate With The Supercomputer Centers In Order To Provide The Best Computing Environment

Support for supercomputer centers was provided without providing support for intermediate equipment, thereby depriving users from doing a substantial amount of their computing on a local basis. NSF Program Managers have not funded mini-supers in the same manner as they did minis and superminis because of the availability of "free" supercomputer time and higher cost of mini-supers.

We need all forms of computing: workstations for productive development and visualization; local super-minis, and mini-supers for the shorter calculations or where programs can be run longer to get the same results; and supercomputers for the exceptional scientific opportunities and "grand challenges" that networked powerful central systems with common programs, data and support provide. To have such an environment requires two components: compatibility and the ability to fund the distributed, small machines.

While I support having the existing centers at the largest, peak power for users whose problems require the maximum computational capabilities, I don't believe they are an "acceptable" way to supply scientific and engineering computation capacity to the larger community. In particular, some of the existing supercomputers in the centers are a poor way to supply computation to certain parts of the engineering community, because most of their small memories, lack of virtual memory, lack of UNIX compatibility (most workstations run UNIX), and network-limited graphics. However, the IBM 3090/600 has both a large physical and virtual memory. Similarly the ETA 10 has large physical and virtual memories and will have a UNIX capability.

## Comparing a Supercomputer Centers Approach With Distributed Minisupers

Enclosed is a table which compares the centers approach with distributed mini-supercomputers which are operated directly by the research community. While the table compares the minisuper with the super, it isn't the intent to eliminate the center. Distributed machines would be managed on a group, department, or university-wide basis in precisely the same fashion as the several thousand minicomputers and super-minicomputers are today. The table compared a single Cray XMP (sans networking) center with a collection of mini-supers costing roughly \$600K (plus interest, but without discount), assuming that most of the operational costs for the mini-super are borne by the user communities institution. In essence, distributing the mini supercomputer transfers the costs of the large central staff, facilities, materials and supplies, travel, networking, etc. from direct NSF funding to the organization using the distributed center. By distributing,



usually one full-time person and a plethora of students maintain a mini and it is expected that this would continue with the mini super environment. Furthermore, many students are now deprived of the operating experience and training of a supercomputer environment.

The cost to NSF that I've used to operate a Center is \$11M, and doesn't include the roughly \$1M required for network access (about \$1K per user), nor the upgrade. The NSF budget amount request for the centers in FY '88 averaged \$9M. The average cooperative agreement request is \$9.7M. The amount needed is somewhat more--\$10.7M. If you include the network which is needed for an overall integrated environment, it gets to roughly \$12M per center. They still must be able to continue at the level of \$2M-\$4M of outside support. In contrast to supercomputers which are increasing in price, superminis are getting cheaper with the large number of new suppliers and approaches, but mini supers have remained in the \$0.5M to \$1.5M price range, probably because of the artificially limited market for these devices.

These cases are provided to compare the distributed and the centralized approach. Each assumes that a user base of 1200 is to be supported which is similar to that of a NSF supercomputer center. Also, it is assumed that a minisuper provided 1/24th the capacity of an XMP/48, or each mini would be equal to 1/6th of processor (1/4-1/3) is probably more realistic. Note both cases favor the distributed approach:

1. 24 superminis, with 50 users each could supply the capacity but not the ultimate capability of the XMP. The NSF cost for 5 years for the computers would be only \$6M versus \$11M. To provide operational assistance to such a large community (50 users) would probably cost NSF an additional 100K, or raise the NSF cost by another \$2.4M. This system provided equivalent capacity of a single center. Certainly, one would not operate 24 superminis in a centers environment!
2. To serve 1200 users, with only 25 users per machine would cost roughly the budgeted FY '88 NSF cost per center. However, each user would have at least twice as much computing capacity, and have no networking limits.

The main environments where supercomputers make sense are the large, centralized, National laboratories which can afford large support staffs, and which need large amounts of computing resources including many shared programs and databases. Also, they require minimal networking.

#### New Machines In Various Price Ranges

A number of new machines that will provide opportunities for service as super and mini supercomputers are described below.

Thinking Machines has introduced a much faster version of the Connection Machine at a price of about \$5M which will be useful for a reasonably large class of problems.

ETA has announced their "Piper" running at 205 speed which is a room cooled computer running Unix.

A large number of conventional and parallel processing computers (multicomputers) exist and are being introduced, all of which offer significant (factors of 2-10) performance and/or performance/price improvements for scientific computing. For example, a new company, Multiflow, based on work NSF funded at Yale on parallel processing just introduced a new \$400K machine which computes at 1/4-1/2 a Cray XMP for Linpack, and higher on the average because it automatically parallelizes over 7 operations per instruction.

Several RISC processors, including MIPS, have introduced computers which have scalar integer performance, characteristic of work done by compilers and operating systems, equal to the Cray XMP.

We can support none of these directly via CISE without trading off some important component of our current activities. The program offers refer computer requests to ASC and have not decided to support mini supers in the same manner that they have supported minicomputers in the past.

#### Center-based Management versus Research-based Management

There is essentially no management of the program based on research needs. The computational science and engineering initiatives we are able to fund are anemic and completely out of balance with the very large centers budget! A real program would be the basis of Wilson's "Grand Challenges", and until each of the disciplines is given the dream and responsibility, they will not deal with the opportunity.

Also, given that essentially all of ASC funding is going into centers we have no funding to understand scientific computing (e.g. benchmarks), to improve productivity through visualization, or new algorithms or new approaches to computation based on parallelism. This lack of understanding will be the first limit of using the next round of supercomputers which are predicated

on a number of processors.

### Bottom Line

Many opportunities exist: The incompatible computing environment at the centers, lack of graphics, novel research results that could only be accomplished with larger resources, performance levels, new user training and population, growth by engineering users, vendor support *outside of IBM's total commitment*, and industrial involvement I am disappointed with the imbalance in the existing program given we are spending so much and cannot address the entire spectrum of requirements. Above all, we have not provided an important style of user-managed distributed computing for our users, including new machines which could provide much more capacity, capability, and training.

### Recommendations

Given the political environment surrounding large centers, I don't believe NSF could withdraw support to any of the centers in order to fund a more balanced computing environment, even though this is what I recommend given that NSF in effect operates a zero sum game. For balance, the funding should be increased to support them at their peak power, along the lines I argued in an earlier memo. At the same time, we need a much stronger CSE Program which is distributed among the research directorates and divisions, along with the computer time. Finally research directorates should encourage users to buy their own smaller, more cost-effective, and in the case of memory, more powerful computers. None of this is happening.

I believe we need both ideas and help in order to have a better balanced program.

GB6

Subject: Alternatives to fund the centers up to their desires and abilities?

To: Director

From: Assistant Director, CISE

Date: 21 August 1987

Based on what I heard at a recent meeting reviewing the centers plans, it is clear that the ASC centers have the ability to absorb an arbitrary amount of funding for the following arguably reasonable things:

- incremental equipment for better balance (e.g. memory, disks, lines)
- upgrades to prevent center obsolescence increased capacity to meet industrial needs and opportunities “batch” and remote visualization equipment for movies
- and lesser priority items including:
- courses
- grand challenges in computational science new technologies and new techniques in parallelism scientists to help in parallelization and visualization interactive visualization at the user level

We cannot possibly meet the requests. The disturbing fact is that NSF is the sole source of support at nearly all of the centers except Illinois, and the degree of support is increasing. Our current approach to funding has literally reduced industrial support. Except for Cornell and Illinois, the centers are really decoupled from industry; they are customers rather than research partners.

I would like to find some other ways to share this incredible funding burden. Here’s my current list of options (ideas):

1. Status quo. NSF funds it all centrally, as we do now in competition with computer science. This is the worst of all possible worlds because the use of the facility is completely decoupled from the supply of the service. By being in CISE, nearly everyone associated with the budget, gets the erroneous conclusion that people working on computer science and engineering research have something to do with the centers. Little or no coupling or use of the centers is made by computer science. The machines aren’t suitable for computing research, nor are adequate funds available for computational science.

If I make the decision to trade-off, it will not favor the centers, but rather centers will be funded at about the same as overall science.

2. Central facility. NSF funds ASC as an NSF central facility. This allows the Director, who has the purview for all facilities and research to make the trade-offs across the foundation.

3. NSF Directorate use taxation. NSF funds it via some combination of the directorates on a taxed basis. The overall budget is set by AD’s. DASC would present the options, and administer the program.

4. Directorate-based centers. The centers (all or in part) are “given” to the research directorates. NCAR provides an excellent model for say BBS, and MPS. Engineering might also operate a facility. I see great economy, increased quality, and effectiveness coming through specialization of programs, databases, and support. This is partially happening.
5. Co-pay. In order to differentially charge for all the upgrades and incrementally nice facilities a tax would be levied on various allocation awards. Such a tax would be nominal (e.g. 5%) in order to deal with the infinite appetite for new hardware and software. This would allow other agencies who use the computer to also help pay.
6. Manufacturer support. Somehow, I don’t see this changing for a long time. A change would require knowing something about the power and throughput of the machines so that manufacturers could compete to provide lower costs. *BTW: Erich Bloch and I visited Cray Research and succeeded in getting their assistance.*
7. Make the centers larger to share support costs. Manufacturers or service providers could contract with the centers to “run” facilities. This would reduce our costs somewhat on a per machine basis.
8. Fewer physical centers. While we could keep the number of centers constant, greater economy of scale would be created by locating machines in a central facility and running them more like LASL and LLNL where each run 8 Crays to share operators, mass storage and other forms of hardware and software support. With decent networks, multiple centers are even less important.
9. Simply have fewer centers. but with perhaps increasing power.
10. Maintain centers at their current or constant core levels for some specified period. Each center would be totally responsible for upgrades, etc. and their own ultimate fate.
11. Free market mechanism. Provide grant money for users to buy time. This might cost more because I sure we get free rides at places like Berkeley, Michigan, Texas and the increasing number of other institutions who do provide megaflops to their users.

I really question how we are going to fund this program in any fashion which permits the facility to be “traded-off” as part of a total research program. Only the disciplines can do this. I believe we should do the following:

1. consolidate equipment in fewer equipment-based centers to reduce cost and operate fewer physical centers at a greater economy of scale
2. have 3 or 4 directorate based centers and 2 or 3 general centers
3. use co-pay as a means to look at real need and as a way to fund specialized facilities such as 35mm movie equipment

Can I have your help on this matter?

MEMORANDUM ...

**DRAFT (NSF Confidential do not reproduce)**

Date: August 25, 1987

To: File

From: Assistant Director, CISE

Subject: Advanced Scientific and Engineering Computing Direction

I'd like to raise several strategic issues for our support of advanced computing for research which go beyond the budgeting for the NSF Supercomputer Centers. This concern is based on these factors:

- inability to create and support a completely compatible distributed
- inability to support new, faster and/or more cost-effective
- the inefficiency of central versus research discipline-based
- inability to support a CSE or a "Grand Challenges" program via the research directorates (which is necessary for a true revolution in

computing environment  
computers on an open  
management

scientific computing)

Background

In the fall, I wrote a memo which was circulated widely within the community (Preparing for Changing Scientific Computing Environment, 30 September 1986). The reactions were bi-modal: "don't rock the boat because at last we have supercomputer time" and "your right, we need a range of compatible computers for a complete environment, furthermore we need support for smaller machines". I am now convinced the hypothesis of the Cray XMP/ VAX-780 gap was correct, and the genesis of the supercomputer program. If the program hadn't been started, then today, researchers would be using Alliant, Convex, DEC, Elexsi, IBM, etc. mini-supers and super-minicomputers, and would have access to significantly more and higher quality computing resources. Independently, Stuart Rice made similar points in a recent speech:

"I have in mind a networked system ... graphics workstations and local supporting intermediate computer and ultimately connects to a supercomputer, with provisions of special devices...

1. Distribution of computer resources distorted by the use of "funny money" ... cash and credit ... Workstations come from grants, supers are "free" ... intermediate machines are indispensable and current funding patterns have to change if ...

2. Dramatic advances in hardware haven't been matched by advances in algorithms and operating systems. ... parallelism is "chicken and egg"

3. The scientific community has become rather inflexible with respect to use of operating systems. ... don't use particular machine features ...

4. ... the scientific community has not been as imaginative as it might in thinking about the uses of computation in research."

#### A Vertically and Horizontally Compatible Computing Environment

Rice's scenario is based on complete horizontal (all the supercomputers) and vertical (supers, mini-supers, and workstation) compatibility. We have total incompatibility across the six centers (including NCAR). They are not compliant with the forthcoming Federal Information Processing Standard, FIPS P1009, Posix. I believe this arcane situation should change, but so far we have no commitment for UNIX at the centers, other than a contract for the ETA 10. Furthermore, NSF is supporting the development of "Unix-like" calls for CTSS, thereby creating a continuing support commitment for a proprietary operating system that I don't believe the government should be funding. The ASC Program simply has been unable to address this problem, thereby limiting the use of machines for entire communities such as VLSI designers and creating more work for all users who now operate with the standard (workstations and minis).

#### Mini-supers Are Required To Operate **With** The Supercomputer Centers In Order To Provide The Best Computing Environment

Aside from the incompatibility that prevents users from moving among the centers, exclusive support for centers has driven support of smaller machines away, thereby depriving users from doing a substantial amount of their computing on a local basis. NSF Program Managers have not funded mini-supers in the same fashion as the traditional mini and super-mini, given the availability of "free" time at supercomputer centers and increased cost of mini-supers.

We need all forms of computing: workstations for productive development and visualization; local super-minis, and mini-supers for 90% of the calculations; and supercomputers for the exceptional scientific opportunities and "grand challenges" that a powerful central system with common programs, data and support would give. To have such an environment requires two components: compatibility and the ability to fund the distributed, small machines.

While I support having several centers at the largest, peak power for the top 5%-10% problems and users, I don't believe they are an acceptable way to supply scientific and engineering computation to 90% of the community. In particular, they are a poor way to supply computation to certain parts of the engineering community, because most of the computers have small memories, lack virtual memory, lack UNIX compatibility (most workstations run UNIX), and lack graphics (limited by the network for the foreseeable future).

#### Comparing a Supercomputer Centers Approach With Distributed Minisupers

Enclosed is a table which shows some of the gains by utilizing mini-supercomputers which are operated directly by the research community. While the table compares the minisuper with the super, it isn't the intent to eliminate the center. Distributed machines would be managed on a group, department, or university-wide basis in precisely the same fashion as the several thousand minicomputers and super-minicomputers are today. The table compares a single Cray XMP (sans networking) center with a collection of mini-supers such as the Alliant FX-8 costing roughly \$600K (sans interest and any discount), assuming that the operational costs for the mini-super are borne by the user community in much the same way the end users fund part of the operational cost of the centers. In essence, distributing the computers eliminates the large central staff and facility expenses of a center. By distributing, usually one full-time person and a plethora of students maintain a mini. Furthermore, many students are now deprived of this operating experience and training.

The cost I've used to operate at Center is \$11M, and doesn't include the roughly \$1M required for network access (or about \$1K per user). The amount in '88 will vary for a center from \$9M (the budget), \$9.7M (agreement), \$10.7M (my request), going to roughly \$11M in '90, but actually requiring more like \$13M to make up for the short fall needed to have modern centers. If you include the network this gets to roughly only \$14M per center, assuming they still are able to generate \$2M-\$4M of outside support. In contrast to supercomputers which are increasing in price, superminis are getting cheaper with the large number of new suppliers and approaches. Note all four cases favor the distributed approach:

1. 16 superminis, with 75 users each could supply the power of the XMP. The cost for the computers would be only \$2M versus \$11M. To provide



operational assistance to such a large community (75 users) would probably cost an additional 100K-200K, or raise the cost to the \$5M range. Such a system could only serve 16 sites conveniently, or 80 sites if they replaced the 5 centers (something that I am **not** advocating).

2. To serve 1200 users, with only 25 users per machine would cost about \$6M. Note, this is roughly equivalent to the user community and number of sites we support today. However, each user would get about 3 times as much computing, and have no networking costs or limits.

3. If NSF provides \$11M per center with only 1200 users, 92 mini-supers could be purchased. Each computer would serve 13 users, each of which would get about 5-1/2 times as much computing power as they do today.

4. If NSF provides \$11M and each machine has only 25 users, 2300 users could be served for the same cost. Each of the old and new users would get about three times the power of existing supercomputer users.

A second case which assumes \$0.1M/year for operations is added to the mini-supers cost, still shows favorable results. The main environments where supercomputers make sense are the large, centralized, National laboratories which can afford large support staffs, and which need large amounts of computing resources including many shared programs and databases. Also, they require minimal networking.

#### New Machines In Various Price Ranges

Thinking Machines is introducing a much faster version of the Connection Machine at a price of under \$10M which will "be the supercomputer" for a reasonably large class of problems.

ETA threatens to supply their "Piper" which is a room cooled computer running Unix. Such a machine should be sold in the "below a \$1M" range.

A large number of conventional and parallel processing computers (multicomputers) exist and are being introduced, all of which offer significant (factors of 2-10) performance and/or performance/price improvements for scientific computing. For example, a new company, Multiflow, based on work NSF funded at Yale on parallel processing is just introducing a new \$400K machine which computes at 1/2 a Cray 1 for Linpack, and higher on the average because it automatically parallelizes over 7 operations per instruction.

We can support none of these either directly via CISE or via the research programs because program officers refer computer requests to ASC. Much of our resources are spent in supporting Cray machines in a mode by which users run small programs on our incompatible operating systems.

#### Center-based Management versus Research-based Management

Our exclusive focus is on highest performance supercomputers and centralized top-down allocation schemes. The allocation overhead for time alone is inefficient and bureaucratic, and is in effect, double jeopardy for researchers. This is uncomfortably unsound now given the "competition and availability" of other non-NSF facilities (roughly 40 IBM 3090's and 4 Crays) and the absence of effective mechanisms for bounding the activities of the centers and evaluating their operation as facilities. I strongly recommend having the various research directorates or divisions allocate the computer time at the centers in order to get the proper focus on managing the complete set of computational resources including computer time, special programs and databases, and a computational science and engineering program (if one ever gets established).

Our exclusive focus on centers, forced by congressional interest, overshadows any real analysis of the evolving needs for research, and prevents us from exploiting obvious advances in place and any unforeseen such as those described above.

#### Computational Science and Engineering Research

There is essentially none! The initiatives we are able to fund are anemic and completely out of balance with the very large centers budget! A real program would be the basis of Wilson's "Grand Challenges", and until each of the disciplines are given the dream and responsibility, they will not deal with the opportunity.

#### Bottom Line

By virtually all measures (the incompatible computing environment at the centers, lack of graphics, novel research results that could only be accomplished with large resources, performance levels, new user training and

population, growth especially by engineering users, vendor support *outside of IBM's total commitment*, and industrial involvement) I am disappointed with the program. Worse, we are spending too much and our future is completely mortgaged to an even costlier future program. Above all, centers have driven out an important style of user-managed distributed computing for 90% of the users, including new machines which could provide much more capacity, capability, and training.

### Recommendations

Given the political environment surrounding large centers, I don't believe NSF would close one of the centers in order to fund a more balanced computing environment, even though this is what I recommend given we are operating in a zero sum game. Thus, the funding should be increased to support them properly, along the lines I argued in an earlier memo. At the same time, we need a much stronger CSE Program which is distributed among the research directorates and divisions, along with the computer time. Finally research directorates should encourage users to buy their own smaller, more cost-effective, and in the case of memory, more powerful computers.

MEMORANDUM

SUBJECT: Advanced Scientific Computing Strategic Opportunities  
TO: FILE  
FROM: Assistant Director/CISE  
DATE: August 31, 1987

I'd like to raise several strategic issues for our support of advanced computing for research which go beyond the budgeting for the NSF Supercomputer Centers. This concern is based on:

o lack of support for distributed minisupercomputer facilities which complement, but do not replace the DASC supercomputer centers

o difficulties and long lead times in the creation and support of a complete compatible distributed computing environment across a range of machines horizontally (supers) and vertically (supers, mainframes, super-minis, mini-supers, and workstations)

o inability to support new, faster and/or more cost-effective computers on an opportunistic basis, (because the centers consume such a large fraction of our resources)

o the inefficiency of central facility versus research discipline-based management

o inability to support a CSE or a "Grand Challenges" program via the research directorates (which is necessary for a true revolution in scientific computing)

o lack of understanding and inability to fund benchmarks dealing with performance and performance/cost issues associated with alternative forms of computing

Background

Last fall, I wrote a memo which was circulated widely within the community (Preparing for Changing Scientific Computing Environment, 30 September 1986). The reactions were bi-modal: "don't rock the boat because at least we have supercomputer time" and "you're right, we need a range of compatible computers for a complete environment, furthermore we need support for smaller machines". I am now convinced the hypothesis of the Cray XMP/VAX-780 gap was correct, and the genesis of the supercomputer program. However, if the program hadn't been started, then today, more researchers would be using mini-supers and super-minicomputers such as Alliant, Convex, DEC, Elexsi, IBM, etc., which would serve many researchers better than the centers are able to. For those researchers who do not need the maximum capability offered by supercomputers, mini supercomputers distributed into the researchers' own environment and connected via networks to the supercomputer centers would offer significantly more and higher quality computing resources to those researchers. Note that capacity refers to throughput and capability refers to turnaround. Independently, Stuart Rice made similar points in a recent speech:

"I have in mind a networked system . . . graphics workstations and local supporting intermediate computer and ultimately connects to a supercomputer, with provisions of special devices . . .

1. Distribution of computer resources distorted by the use of "funny money" . . . cash and credit . . . Workstations come from grants, supers are "free" . . . intermediate machines are indispensable and current funding patterns have to change if . . .
2. Dramatic advances in hardware haven't been matched by advances in algorithms and operating systems. . . . parallelism is "chicken and egg"
3. The scientific community has become rather inflexible with respect to use of operating systems. . . . don't use particular machine features . . .
4. . . . the scientific community has not been as imaginative as it might in thinking about the uses of computation in research."

### A Vertically and Horizontally Compatible Computing Environment

Rice's scenario is based on complete horizontal (all the supercomputers) and vertical (supers, mini-supers, and workstation) compatibility. We now have six incompatible centers (including NCAR). Since UNIX is truly becoming an important standard under

the forthcoming Federal Information Processing Standard, FIPS P1009, Posix, I believe this situation must change. Unfortunately, NSF at SDSC is supporting the development of "Unix-like" calls for CTSS, thereby creating a continuing support commitment for a proprietary operating system. The ASC Program has been slowly addressing this problem. Until UNIX is in place the existing situation limits the use of the centers' supercomputers for entire communities such as VLSI designers and creating more work for their users who now operate with the standard (workstations and minis).

### Mini-supers Are Required to Operate **With** The Supercomputer Centers In Order To Provide The Best Computing Environment

Support for supercomputer centers was provided without providing support for intermediate equipment, thereby depriving users from doing a substantial amount of their computing on a local basis. NSF Program Managers have not funded mini-supers in the same manner as they did minis and superminis because of the availability of "free" supercomputer time and higher cost of mini-supers.

We need all forms of computing: workstations for productive development and visualization; local super-minis, and mini-supers for the shorter calculations or where programs can be run longer to get the same results; and supercomputers for the exceptional scientific opportunities and "grand challenges" that networked powerful central systems with common programs, data and support provide. To have such an environment requires two components: compatibility and the ability to fund the distributed, small machines.

While I support having the existing centers at the largest, peak power for users whose problems require the maximum computational capabilities, I don't believe they are an "acceptable" way to supply scientific and engineering computation capacity to the larger community. In particular, some of the existing supercomputers in the centers are a poor way to supply computation to certain parts of the engineering community, because most of their small memories, lack of virtual memory, lack of UNIX compatibility (most workstations run UNIX), and network-limited graphics. However, the IBM 3090/600 has both a large physical and virtual memory. Similarly the ETA 10 has large physical and virtual memories and will have a UNIX capability.

### Comparing a Supercomputer Centers Approach With Distributed Minisupers

Enclosed is a table which compares the centers approach with distributed mini-supercomputers which are operated directly by the research community. While the table compares the minisuper with the super, it isn't the intent to eliminate the center. Distributed machines would be managed on a group, department, or university-wide basis in precisely the same fashion as the several thousand minicomputers and super-minicomputers are today. The table compared a single Cray XMP (sans networking) center with a collection of mini-supers costing roughly \$600K (plus interest, but without discount), assuming that most of the operational costs for the mini-super are borne by the user communities institution. In essence, distributing the mini supercomputer transfers the costs of the large central staff, facilities, materials and supplies, travel, networking, etc. from direct NSF funding to the organization using the distributed center. By distributing, usually one full-time person and a plethora of students maintain a mini and it is expected that this would continue with the mini super environment. Furthermore, many students are now deprived of the operating experience and training of a supercomputer environment.

The cost to NSF that I've used to operate a Center is \$11M, and doesn't include the roughly \$1M required for network access (about \$1K per user), nor the upgrade. The NSF budget amount request for the centers in FY '88 averaged \$9M. The average cooperative agreement request is \$9.7M. The amount needed is somewhat more--\$10.7M. If you include the network which is needed for an overall integrated environment, it gets to roughly \$12M per center. They still must be able to continue at the level of \$2M-\$4M of outside support. In contrast to supercomputers which are increasing in price, superminis are getting cheaper with the large number of new suppliers and approaches, but mini supers have remained in the \$0.5M to \$1.5M price range, probably because of the artificially limited market for these devices.

These cases are provided to compare the distributed and the centralized approach. Each assumes that a user base of 1200 is to be supported which is similar to that of a NSF supercomputer center. Also, it is assumed that a minisuper provided 1/24th the capacity of an XMP/48, or each mini would be equal to 1/6th of processor (1/4-1/3) is probably more realistic. Note both cases favor the distributed approach:

1. 24 superminis, with 50 users each could supply the capacity but not the ultimate capability of the XMP. The NSF cost for 5 years for the computers would be

only \$6M versus \$11M. To provide operational assistance to such a large community (50 users) would probably cost NSF an additional 100K , or raise the NSF cost by another \$2.4M. This system provided equivalent capacity of a single center. Certainly, one would not operate 24 superminis in a centers environment!

2. To serve 1200 users, with only 25 users per machine would cost roughly the budgeted FY '88 NSF cost per center. However, each user would have at least twice as much computing capacity, and have no networking limits.

The main environments where supercomputers make sense are the large, centralized, National laboratories which can afford large support staffs, and which need large amounts of computing resources including many shared programs and databases. Also, they require minimal networking.

#### New Machines In Various Price Ranges

A number of new machines that will provide opportunities for service as super and mini supercomputers are described below.

Thinking Machines has introduced a much faster version of the Connection Machine at a price of about \$5M which will be useful for a reasonably large class of problems.

ETA has announced their "Piper" running at 205 speed which is a room cooled computer running Unix.

A large number of conventional and parallel processing computers (multicomputers) exist and are being introduced, all of which offer significant (factors of 2-10) performance and/or performance/price improvements for scientific computing. For example, a new company, Multiflow, based on work NSF funded at Yale on parallel procesisng just introduced a new \$400K machine which computes at 1/4-1/2 a Cray XMP for Linpack, and higher on the average because it automatically parallelizes over 7 operations per instruction.

Several RISC processors, includingMIPS, have introduced computers which have scalar integer performance, characteristic of work done by compilers and operating systems, equal to the Cray XMP.



We can support none of these directly via CISE without trading off some important component of our current activities. The program offers refer computer requests to ASC and have not decided to support mini supers in the same manner that they have supported minicomputers in the past.

### Center-based Management versus Research-based Management

There is essentially no management of the program based on research needs. The computational science and engineering initiatives we are able to fund are anemic and completely out of balance with the very large centers budget! A real program would be the basis of Wilson's "Grand Challenges", and until each of the disciplines is given the dream and responsibility, they will not deal with the opportunity.

Also, given that essentially all of ASC funding is going into centers we have no funding to understand scientific computing (e.g. benchmarks), to improve productivity through visualization, or new algorithms or new approaches to computation based on parallelism. This lack of understanding will be the first limit of using the next round of supercomputers which are predicated on a number of processors.

### Bottom Line

Many opportunities exist: The incompatible computing environment at the centers, lack of graphics, novel research results that could only be accomplished with larger resources, performance levels, new user training and population, growth by engineering users, vendor support *outside of IBM's total commitment*, and industrial involvement I am disappointed with the imbalance in the existing program given we are spending so much and cannot address the entire spectrum of requirements. Above all, we have not provided an important style of user-managed distributed computing for our users, including new machines which could provide much more capacity, capability, and training.

### Recommendations

Given the political environment surrounding large centers, I don't believe NSF could withdraw support to any of the centers in order to fund a more balanced computing

environment, even though this is what I recommend given that NSF in effect operates a zero sum game. For balance, the funding should be increased to support them at their peak power, along the lines I argued in an earlier memo. At the same time, we need a much stronger CSE Program which is distributed among the research directorates and divisions, along with the computer time. Finally research directorates should encourage users to buy their own smaller, more cost-effective, and in the case of memory, more powerful computers. None of this is happening.

I believe we need both ideas and help in order to have a better balanced program.

GB6

Issues to address in the DASC Long Range Plan

What is the New Technologies charter?

What is DASC's role vis a vis Computational Science and Engineering (CSE) in encouraging the disciplines to effectively organize and utilize supercomputers?

How can the users build quality large codes that have traditionally come out of the efforts of researchers and professionals built over 5-10 years?

The centers have to demo that they are doing something of added value that could not better be done in a single center, given the expensive NSFnet.

The network cost-effectiveness

Subject: Alternatives to fund the centers up to their desires and abilities?

To: Director

From: Assistant Director, CISE

Date: 21 August 1987

Based on the what I heard at a recent meeting reviewing the centers plans, it is clear that the ASC centers have the ability to absorb an arbitrary amount of funding for the following arguably reasonable things:

- incremental equipment for better balance (e.g. memory, disks, lines)

- upgrades to prevent center obsolescence

- increased capacity to meet industrial needs and opportunities

- "batch" and remote visualization equipment for movies

and lesser priority items including:

- courses

- grand challenges in computational science

- new technologies and new techniques in parallelism

- scientists to help in parallelization and visualization

- interatctive visualization at the user level

We cannot possibly meet the requests. The disturbing fact is that **NSF is the sole source of support** at nearly all of the centers except Illinois, and the degree of support is increasing. Our current approach to funding has literally reduced industrial support. Except for Cornell and Illinois, the centers are really decoupled from industry; they are customers rather than research partners.

I would like to find some other ways to share this incredible funding burden. Here's my current list of options (ideas):

- 1. Status quo.** NSF funds it all centrally, as we do now in competition with computer science. This is the worst of all possible worlds because the use of the facility is completely decoupled from the supply

of the service. By being in CISE, nearly everyone associated with the budget, gets the **erroneous** conclusion that people working on computer science and engineering research have something to do with the centers. Little or no coupling or use of the centers is made by computer science. The machines aren't suitable for computing research, nor are adequate funds available for computational science.

If I make the decision to trade-off, it will **not** favor the centers, but rather centers will be funded at about the same as overall science.

**2. Central facility.** NSF funds ASC as an NSF central facility. This allows the Director, who has the purview for all facilities and research to make the trade-offs across the foundation.

**3. NSF Directorate use taxation.** NSF funds it via some combination of the directorates on a taxed basis. The overall budget is set by AD's. DASC would present the options, and administer the program.

**4. Directorate-based centers.** The centers (all or in part) are "given" to the research directorates. NCAR provides an excellent model for say BBS, and MPS. Engineering might also operate a facility. I see great economy, increased quality, and effectiveness coming through specialization of programs, databases, and support. This is partially happening.

**5. Co-pay.** In order to differentially charge for all the upgrades and incrementally nice facilities a tax would be levied on various allocation awards. Such a tax would be nominal (e.g. 5%) in order to deal with the infinite appetite for new hardware and software. This would allow other agencies who use the computer to also help pay.

**6. Manufacturer support.** Somehow, I don't see this changing for a long time. A change would require knowing something about the power and throughput of the machines so that manufacturers could compete to provide lower costs.

**7. Make the centers larger to share support costs.** Manufacturers or service providers could contract with the centers to "run" facilities. This would reduce our costs somewhat on a per machine basis.

**8. Fewer physical centers.** While we could keep the number of centers constant, greater economy of scale would be created by locating machines in a central facility and running them more like LASL and LLNL where each run 8 Crays to share operators, mass storage and other forms of hardware and software support. With decent networks, multiple centers are even less important.

**9. Simply have fewer centers.** but with perhaps increasing power.

**10. Maintain centers at their current or constant core levels for some specified period.** Each center would be totally responsible for upgrades, etc. and their own ultimate fate.

**11. Free market mechanism.** Provide grant money for users to buy time. This might cost more because I'm sure we get free rides at places like Berkeley, Michigan, Texas and the increasing number of other institutions who do provide megaflops to their users.

I really question how we are going to fund this program in any fashion which permits the facility to be "traded-off" as part of a total research program. Only the disciplines can do this. I believe we should do the following:

1. consolidate equipment in fewer equipment-based centers to reduce cost and operate fewer physical centers at a greater economy of scale
2. have 3 or 4 directorate based centers and 2 or 3 general centers
3. use co-pay as a means to look at real need and as a way to fund specialized facilities such as 35mm movie equipment

Can I have your help on this matter?

# Advanced Scientific Computing Program Evolution

## Original OASC Program Goals

- supply supercomputer power to scientific and engineering community
- train scientific and engineering community to use supercomputers
- support U. S. supercomputer industry
- provide for remote access by users
- stimulate development of a "rich and powerful" scientific/engineering computing environment **for all** as measured by total power and utility*

## Evolving Centers Program

- provide the highest performance service with the latest supercomputers, acting as "beta" sites for the manufacturers
- provide generic applications software libraries
- maintain "consortia nets", and connect to NSFnet and regional nets for user access
- train "selected or strategic" users in various disciplines (30/center)
- allocate time; report use by disciplines, user experience, user job size, and geography; review effectiveness for users; and review final scientific/engineering output
- co-ordinate and establish a common program, user interface, and graphics environment across centers to maximize user effectiveness ?*
- establish key links with hardware and software suppliers?*
- establish a program for industrial users?*

## CISE-wide Programs in CSE, New Technologies, and DCCR

- benchmark and understand performance and cost-effectiveness of various conventional scientific, multi-vector processor computers
- enhance training for all disciplines to include vector processing
- provide advanced graphics techniques (visualization)
- provide parallel processing environments and encourage applications appropriate to efficient use of multiprocessor supercomputers
- research on advanced algorithm and software across disciplines
- research new technologies for potential performance breakthroughs

## Other NSF CSE Programs

- develop and maintain area(x)-specific, common user communities for common programs and datasets
- sponsor and organize the "grand challenges" in computational science

## Changing Networking Activities-NCRI

- Originally provided network to link centers to users, to each other, and to other networks.
- Currently links to other databases, going beyond traditional computing
- Extend network to support entire research infrastructure*

Gordon Bell; 10 March 1987

### **Why was it necessary to establish a central directorate for computing?**

CISE recognizes the pervasiveness of the computer in society today and the unique opportunities for computing at this time.

The information society, which is the largest sector of the economy is based on computing and communication. Just as mechanisms were the basis of the industrial revolution, computers are the basis of the information revolution.

Computing is found in virtually every scientific and engineering discipline as a base either as a tool or a component, and it is a science in its own right. In science, the Nobel Laureate, Ken Wilson, and head of the Cornell Theory Center housing one of the NSF supercomputer centers, explains computation as the third paradigm of science. The first being theory, and second, experimentation.

History has shown that government funding of computing research has been the main driving force of the revolution in computing that has become the largest industry today.

Finally, today, we have a new opportunity vis a vis parallelism to come off the technology evolutionary path of the last few decades that provide only x10 of performance per decade.

### **How Do you see Computing research affecting Competitiveness?**

Directly through products. We have history of revolution that has occurred by funding university research -The Army funded Eniac and Edvac at Penn, the first computers that became the basis of modern computers, and the designers went on to create Univac. At MIT, ONR funded Whirlwind, from which came core memories, real time, air defense, air traffic control computers, interactive computing and the first computer aided manufacture. Digital Equipment Corporation came almost directly from the Whirlwind effort and team. Timesharing was first implemented at MIT; this became the basis of all modern computing. Graphics research, initially at the University of Utah, became the basis of all workstations and PC's, its how computers are beginning to be truly useful to everyone. ARPA funded communications networks for computing. The artificial intelligence-based expert system at Digital to specify how computers are put together was first prototyped at Carnegie Mellon U. This was the basis of the emerging AI industry. Universities are the main source of ideas and programs in VLSI design.

Only this week, these example of NSF funded projects came across my desk: Don Knuth's program, TEX, is now the basis for modern typesetting of scientific and mathematics manuscripts. Two different parallel processing schemes, are now implemented by 5 companies. A research at Utah has just implemented a text searching scheme that promises to be able to retrieve any text in any size database in virtually 0 time. Kamakaha's algorithm at BTL came out of extension of his thesis work at UC/B.

The supercomputer centers produce results regularly: America's cup, Kodak Material, Corning (1/6 throw away) simulation of the new superconducting materials, search for cold virus serum, molecular modeling and computational chemistry, we even have work to use the computer as a computational telescope.

Computers are critical to CAD,CAM,CAI, CAI, ... in every environment from home, office, laboratory, vehicle, or factory. We especially are focusing on hardware in this budget, note the increase in the MIPS area.

Finally, we still have a + balance of trade in computers, but its fading fast. Japan is breathing down our neck in every phase of R...D, and every area from AI to payroll.

Bottom line: We have no trouble in measuring results, including gestation times. It is quite rapid, and it can and must be even faster.

### **What are you doing to help the education process?**

I mainly believe that the big force that drives the education process comes from the right balance between research and teaching by first rate researchers. Much of research comes from student questions.

Let me give a homely example of the interaction of teaching and research. I took 6 years off from Digital to teach and do research at Carnegie Mellon university from 66-72. I wanted to explain how simple computers were and to have computers design them. We came up with 2 notations, that later became languages to describe, simulate and ultimately now to begin to automatically design computers. The text we wrote is still a classic on computers, and many simulators use the language, and at least one company sells the program. All of this came out of a research direction and drive that was largely pedagogical.

I also believe that the work we are doing indirectly in CAI in some of the leading universities will ultimately filter into all forms of education.

### **Are you familiar with Rep Sabo's Proposal to have NSF fund the Phase I centers by cutting 15% from the Phase II budgets, and then ultimately go to a free market for all supercomputer service?**

The Phase I program was established to buy computer time from various organizations, including three companies and three universities (Colo state, minn, Purdue) who had supers. We had no long term commitment for support and the contract was clear from the beginning that we were not going to continue support when our own phase II, centers became operational. The phase II centers are all now operational, pretty much according to our plan. By cutting our Phase II centers back 15%, would be a disaster; we simply can not maintain the systems at the performance levels we need, that is having the latest, and highest performance computer available on the market. This requires amortizing a computer over 3.5-5 years, the gestation time for a supercomputer. I do not support the concept of a "free market mechanism" for



machine time at this time, whereby anyone can supply cycles. This mechanism didn't work and was the main drive why the government had to step in and form the ASC program in the first place.

I am in the process of reviewing whether we have adequate funds to maintain our existing centers with the latest computers. It looks as if we are going to have to need more funds. I am not requesting more at this time, but believe we want significantly more help from computer suppliers, several of the states, some of the universities that host the computers, and industrial users. I believe the government is paying too much of the freight.

### **Are you happy with the Program?**

Yes. We have 6K people on 2K projects, at 200 sites in all states. We see exciting results almost daily.

One of the great benefits from the program to date is the side effect of causing a number of great universities to acquire their own supercomputer. I don't believe any great university can afford not to have this kind of capability. For example, Berkeley has a small Cray XMP and an IBM 3090/200, Texas, Ohio State, Minnesota have or acquiring Crays. I hope the ETA computer will be successful, and replace the CDC 205's at various universities that have them. Michigan got the second 3090 after the center at Cornell. IBM has installed 40-3090/200's which could supply significant computing power (each processor of a dual is about equal to a Cray 1). In fact today, I estimate that we have the equivalent of over 110 times the Cray 1 available to the university research community, about 70 of this is in unis, about 40 at the centers (including NCAR) in 4 Cray xmps, 1-Cray 1, 1-3030/400, and 2-CDC 205's that are to be replaced with a machine of 20x a Cray 1.

As an alternative, three companies are building and installing mini-supers, all of which can do many of the tasks supers can do on a cost-effective basis. Many more designs are in the wings.

As the person responsible for getting about 3000 computers of the minicomputer price class into the scientific community in the form of VAX, I think the future will give us lots of options in the way to do computing. Today, supercomputer users generally access supers at the end of a very slow network. This limits their own abilities in a different way, particularly in being able to visualize results. Many things (other centers, superminis, and networks) have changed since the establishment of the centers program, and we must continually evaluate the options for the future.

In all scenarios, I continue to see the need for a few centers which have the latest and fastest computers.

Subject: Congressional Review of the Advanced Scientific Computer Program

To: Erich Bloch  
Ray Bye  
Charles Brownstein  
Mary Clutter  
John Connolly  
James McCullough  
John Moore  
Paul Rotar  
Steve Wolff

From: AD/CISE (Gordon Bell)  
Date: 7 November 1986

The official review to be completed in early December when we present to the House Science and Technology Committee calls for:

"NSF to submit a long-range plan, by December 1986, for implementing:

- development and review of resource allocation policies;
- assessment of supercomputer manpower and training needs;
- evaluation of Phase I and Phase II national supercomputer centers;
- development and maintenance of national supercomputer centers and resource networks;
- and research needs and support for computational mathematics and the development of software, algorithms, and network technologies."

In the internal assessment of the program by the Program Evaluation Staff, I think we want to collect and look at the following data:

I. Complete inventory of scientific and engineering computers at universities (domestic and foreign) including supers/large mainframes (CDC, Cray, IBM 3080/3090); FPS array processors; mini-supers (Alliant, Convex, SCS)

... for example Convex has installed 15 machines, at 1/3-1/5 a Cray XMP processor; super-minis ( $\geq$ VAX 8600, 4381-12); and other  $> 1$  Mflop machines... if they are significant. The criteria is roughly 1 Mflop Linpack, which allows a machine to deliver the equivalent of one hour of Cray time per day.

II. Cost, service, training, and quality of our centers, including comparisons with commercial (eg. Boeing), Phase I, and other centers (eg. U. of Texas).

III. Use and training versus time at the centers by: discipline, geography (including the remoteness), university, and experience as a supercomputer user. The analysis would be carried out both on an individual and project basis. The goal is to determine "fairness" and penetration of use into the scientific and engineering disciplines.

IV. Scientific and engineering output. This is hard to evaluate. The amount of code and common databases running at the centers as community programs might be one measure. We should look at the various classes of code and see how effective the computation is at solving a problem. For example, we need estimates on the computation times in the various disciplines.

Our plan format should be a tone page line of capabilities at various phases:

supply of services; training and user base; networking; graphics; effect on standards; parallelism; establishment of new scientific and engineering communities with particular community databases and programs; leverage of additional resources

Dr. Jim Decker  
Department of Energy

Dear Jim:

I gave a talk on parallel processing at Argonne yesterday at a two week summer session sponsored by DOE and NSF which stimulated several thoughts.

1. More formal interaction between DOE Labs and NSF (or the universities). Dave Nelson was going to sponsor a get together. We still need to try to see if a more formal interaction plan could be beneficial in regard to technology transfer with the universities.

2. Alternatives for the teraflop computer. Having looked at the alternatives to getting a teraflop, the only one which looks doable in a short time frame is a large, simd approach as used in the Connection Machine. With the next generation of CMOS, I believe such a machine could be built which only needs to be a factor of 50-100 times larger than the current Connection Machine. A factor of 2-8 in clock and 8-32 in size looks doable. The big hitch in this is the ease of use. My own belief is that it is much more programmable than the approach you outlined in your approach to tie 1000, 1 gigaflop computers together in a hypercube... here we have a pretty good idea that the programming is very hard, except for problems involving small program kernels.

Given the lab's recent history of supporting nearly all of the unsuccessful computers starting with Star and including HEP, several toy hypercubes (Intel, T-series), S1, PUP, and Berkeley's mP, I can't understand why you are not looking at the single approach that will supply supercomputer power today, can be programmed, and can be extended.

The bottom line on this is to urge you to take an active role in understanding the Connection Machine, because I think it is the only idea that will work for your time-consuming codes. Jack Dongarra would like to have a Connection Machine, but it is out of his budget range. Also, I believe it needs to be at a lab where it can be put to work now on real problems.

3. I learned that Argonne is getting a uniprocessor Cray. Can I urge you to consider buying time either at the University of Illinois or at one of your own centers and not supporting such fiscal irresponsibility. The one thing I now understand about supercomputer centers is that you want as few as possible and that each one needs to be as large as possible in order to provide the best support which is also at the lowest cost. A uniprocessor Cray at an isolated site is probably the most expensive, and poorest form of computing that one can obtain.

At any rate you have the above gratuitous advice.

I enjoyed the interaction with the Argonne group and hope they have some impact and utility within your department and continue to believe they are an important component of the Computer Science Research community.

Sincerely,

Gordon Bell  
Assistant Director  
(and taxpayer)

bcc:  
E Bloch

**The Future Direction of  
High Performance, Scientific and Engineering Computing  
Gordon Bell, 408-732-0400  
Ardent Computer Corporation; Sunnyvale, California 94086**

**Keywords**

future supercomputing, high performance computation, performance/price, parallelism, politics and policy, computational science

**Summary**

Many high performance computers (such as supercomputers) are emerging from the 70 existing and newly formed companies to exploit gains in the technologies of automatic parallelization of programs, reduced instruction-set computers, integrated circuits, and algorithms. A recent report from the Office of Science and Technology Policy has urged the adoption of a High Performance Computing effort. By 1995 a computer capable of  $10^{12}$  operations per second should be possible, providing a factor of almost 1000 gain over today's supercomputer. Dis-economies of scale suggest that a distributed approach is also needed. To reach a teraflop will require a concerted effort involving computer engineers, computer scientists, and computational scientists that to date is lacking.

**Introduction**

Spurred by a number of innovations from both the industrial and academic research establishments made possible by VLSI and parallelism, we can expect the next generation of scientific and engineering computing to be even more diverse and exciting than all others. The research accomplishments have been stimulated by DARPA's Strategic Computing Initiative (SCI), and a review of this activity prompted this paper. However, without the scientific base, creative talent, and infrastructure to design complex VLSI chips, the innovative machines would not be possible. A variety of established and newly formed companies have organized to exploit the new technology. Table 1 gives the number of companies building high performance computers for scientific and engineering applications.

Table 1. Companies building high performance computers.

Kind of Computer	On Market	Developing	Dead	Recent*
Supercomputers	5	2	1	5
Vector Mainframes	4	?	1	1
Mini-supers	6	2	1	8
Graphic Supers	2	2	0	2
Total Supers	17	6	3	16
Array Processors	9	0	4	7
Massive Data Parallel	3	1	0	3
Multiprocessors	9	2	1	11
Multicomputers	15	0	0	15
Total (including supers)	53	9	8	52

Superminis	6	0	2	2
RISC-based computers	5	0	1	6

\* effort started since 1983

The impressive gains in performance and large number of new companies demonstrate the effectiveness of the university-based research establishment and their ability to transfer technology to both established companies and start-up firms. If the current program is followed, it appears that a computer capable of teraflop execution could be constructed by 1995. Figure 1 and 2 provide a taxonomy and projection of growth in performance for various computers.

Three kinds of computers are emerging, segmented according to "general purposeness":

- General purpose computers are multiprogrammed and time-shared, depending on their price, and can handle a variety of applications at a given time. General purpose computers include supers, mainframes, various minis, and workstations.
- Run-time defined, applications specific computers are used on only a few problems, on a one-at-a-time basis. These mono-programmed computers are useful for a more limited class of problems where a high degree of data parallelism exists. This class of computers can achieve a factor of 10-100 increase in some combination of performance and performance/price.
- Applications-specific computers which only solve one problem such as speech or image processing. By binding the application in hardware and software at design and construction time, a factor of 100-10,000 increase in some combination of performance or performance/price is possible.

The following sections describe the innovations in:

- "main line" scientific and engineering computing using vector multiprocessing and following the archetypical "Cray",
- new, very fast, Plain Old one chip Processors (POPs) implemented as uni-microprocessors,
- multiprocessors and multicomputers that result in using POPs,
- more unconventional, highly parallel machines that are beginning to enter the market, and
- highly parallel machines in the research phase to watch.

Several programming models are presented which cover the vast array of machines. The computer space is given in the simplified taxonomy, Fig. 1.

Two possible computer structures could lead to a computer capable of teraflop operation. For special situations, we would expect to see a number of machines executing at the tera-op rate by 1995.

The final sections describe the implications for use, together with actions required by government agencies, users, and the computer science and engineering research community.

**Cray-influenced Main Line of "supers, mini-supers, graphics supers, and to**

### *emerge, personal supers"*

The "main line" of scientific and engineering "super" computer development follows the Cray formula: having the fastest clock; simple, pipelined scalar instruction-set; a very complex pipelined vector processing section; and multiprocessors. Machines have formed in each (price) class: the "Cray" or supercomputer at \$10-20M, the "Crayette" or minisuper at under \$1M, and the emerging graphics super at around \$100K. Within two years a personal super costing less than \$50,000 is possible. A super formed from a one chip vector processor as an "add-on" for a personal computer will also appear in the early 90's. For the purpose of this paper we define the "x"-super (where x = {\_, mini, graphics, personal, micro}) contains as many vector processors and runs at the highest possible clock to still meet the class (price) constraint.

The Japanese computer industry has aimed at having the fastest uniprocessor supercomputers as demonstrated in NEC's SX-2 and Hitachi's S-820/80. These uniprocessors provide the fastest single stream execution, but Crays still have more aggregate throughput. Though its parallelizing compiler, the Cray YMP is now the world's fastest supercomputer (Dongarra, 1988) and will remain so until the Japanese start building multiprocessors.

The next Cray after the Cray YMP (8 Processors, each performing at 1.2-1.5 times the XMP processor) and the Cray 3 (16 processors at 2X the Cray 2) is the Cray 4. The Cray 4, to be available in '92, operates at 1 ns clock rate and delivers 128 Gflops, using 64 processors. Figure 3 shows the evolution of clock speed, number of processors, and aggregate computing power for "Cray" computers over a thirty year period beginning with the CDC 6600.

In order to get this amount of power on a single program will require some form of parallel programming by the user. For example, the Long Range Global Weather Modelling program at NCAR was manually parallelized and operates at over 400 Megaflops on the Cray 416. Given the relatively small memory and the growth in application program size, users are beginning to utilize the XMP in parallel mode to fully utilize its processors.

There appears to be no Economy of scale across the supercomputer classes. In fact, lower cost, high volume products result in an inherent dis-economy of scale in performance/price as we examine below in a comparison between the Cray YMP and Ardent's Titan Graphic Supercomputer. ETA's small scale supercomputer costs about \$1 million and provides over 30 megaflops, or 30 flops/s/\$-- which is 50% better than its high performance, liquid nitrogen cooled 8 processor ETA 10.

### **Plain Old Processors (POPs) Fuel High Performance Computers**

The future of the uniprocessor is a simple, risc-based architecture enabling high speed execution and low cost, one chip (i.e. microprocessor) implementation. By adding an attached vector processor, users can see a very bright picture for scientific and engineering computation in workstations and as simple computers. Similarly, by utilizing ECL gate arrays, it is relatively easy to build processors which operate at 200 Mhz (5 ns clock) and be available by 1990. One such company, Key Computer is doing just this for a minisuper priced uni- and multi-processor. Prisma is building a GaAs



based computer using the SUN Sparc architecture.

The projected evolution of the leading edge one chip POPs is given in Table 2. Unlike the leading edge clock, POP clock speed has evolved a factor of 5 in four years or at a factor of 1.7 per year since the processor is entirely on a chip. Shifting to ECL will give an aggregate speed-up of a factor of 20 over a 6 year period, or a factor of 1.8 per year.

Table 2. Past, presented and project clock and speed for one chip processors (POPS)

time	clock (mhz)	PkMips	Mflops	Mflops vector unit
'86	8	5	1	-
'87	16	10	2	16
'88	25	16	5	25
'89	40	25		
ECL shift				
'90	80	50	10	100
'92	160	100	20	200

### **Computers That Come From Low-cost, and Fast Micros**

The very fast CMOS and soon to come ECL microprocessor will push every computer and be a good component for both multiprocessors and multicomputers. Also, the micros can be used in a redundant fashion like Stratus has pioneered to increase reliability and build what is fundamentally a hardware fault-free computer.

Multiprocessors. The next generation, high performance micros are *at last* designed for building "multi's" (Bell, 1985) multiple microprocessor computers. Several thousand multi's are now in operation and I believe this will become the mainline for traditional shared computers and smaller workstations. However, given the speed and simplicity of POPs, users may ask: why bother with so much performance? By 1990 workstations with 5-10, 20 mips processors attached to a shared bus in a "multi" configuration and selling for under \$50K will exist. However, the 50 mips microprocessor will place much pressure on the viability of the multiprocessor.

The utility of the multiprocessor as a general purpose device is proven because it can be used in a multiprogrammed and time-shared fashion. It is also the object of training and research in parallel processing by the computer science community because it provides the most general purpose tool in that it can provide an environment for a number of computational models.

The Alliant and Convex mini-supercomputers, and Ardent and Stellar graphics supercomputers have vector processors, for general purpose scientific and engineering computation. Other approaches to large multiprocessors do not have vector facilities, and hence may not be viable or performance/price competitive since automatic compilation of parallel constructs to utilize a large number of scalar processors hasn't taken place. Furthermore, it is difficult to build very large multiprocessors as cheaply as the multicomputers described below in order to get power comparable to a supercomputer. Hence, multicomputers have been demonstrated and have gained utility for user-explicit, parallel processing of single problems at supercomputer speeds.

Given the ease of building very high performance multiprocessors based on the above POPs, a revolution is likely in the traditional mainframe and minicomputer industry that sell "code museums" for existing programs. Using this multiprocessor approach, a variety of computers which execute programs at the 100-400 million instruction per second rate and cost around \$500,000 will soon enter the market. Such a structure provides a factor of 2-4 times the power of IBM's largest mainframe at 1/20th the cost.

Hypercubes and succeeding generation multicomputers. Seitz and his colleagues at Cal Tech developed a large, multicomputer structure known as the hypercube in the early 80's. By using commodity micros and interconnecting them in a hypercube or grid (the Inmos Transputer) each computer can pass messages to one another. Today's multicomputer provide substantially higher performance for message passing than their first generation ancestors. For a particular application, a factor of 10 in price/performance over the main line supercomputers has been observed.

Multicomputers are not generally applicable to all problems and are usually mono-programmed since they are only used on one program at a time. Hypercubes now exist with 32-1024 computers, being manufactured by about a half dozen companies, and several hundred are in use. Programs have to be rewritten to utilize the multicomputer message passing system, but the peak performance (see Table 3) and price performance appears to be worth the effort, as a lab can have its own "Cray" for a particular problem.

Table 3. Multicomputer generations, Seitz (1988).

	First 83-87	Second 88-92	Third 93-97
Nodes			
MIPS	1	10	100
Mflops scalar	0.1	2	40
Mflops vector	10	40	200
Memory (Mbytes)	0.5	4	32
No. of nodes*	64	256	1024
Message time (us)**	2K	5	0.5

\*Typical system. Maximum system is roughly four times larger.

\*\*100 Mbyte packet

The European multi-computer counterpart to the hypercomputer is based on Inmos's Transputer. A transputer is a processor with ?? bytes of on chip memory and four, full duplex interconnection ports operating at 20 megabits/sec, for passing messages to other Transputers. Several companies are building general purpose, multi-computers by connecting a large number of transputers together. The transputer is proving especially useful system to build application-specific systems for everything from communications to robots.

### **Supercomputers: price, performance, parallelism, and politics**

Given the incredibly strong lobby for supercomputing in government, academe, and industry it is worth looking at balancing computing away from strong centrality and to a fully distributed approach. Just as distributed computing using minicomputers, workstations, and personal computers has a strong role in computing compared with mainframes, fully distributed supercomputing is likely to take on a similar role with respect to traditional supercomputers. This section looks at the relationship of price, performance, and the politics (and prestige) of two approaches to computing. Table 4 gives the purchase price, performance, performance/price for several benchmarks run both sequentially and in parallel for a Cray YMP and Titan graphics supercomputer.

Observe that for the purchase price of the YMP, one could have 166 graphics supercomputers in a highly distributed fashion for personal, project, or departmental use. The simple model ignores operating costs, which in the case of a central computer are quite visible. In the distributed approach operations cost (e.g. file backup) are buried in the organization. Similarly the cost of support, purchasing and maintaining software, and maintaining files may vary using the two approaches.

Obviously, all of the benchmarks run longer on the slower machine. The stretch factor is the increased time that a program runs using the Titan as compared to runtime on the Cray YMP. Also, associated with each benchmark is the cost-effectiveness or performance/price (e.g. Megaflops/sec/\$) of YMP versus Titan approach.

Performance and Cost-effectiveness. The range of results compare with an analysis of Titan and the Cray XMP by Misaki and Lue at the Institute for Supercomputer Research (1988) where for scalar and vector loops, the Cray was roughly a factor of five and ten

faster respectively. The Whetstone benchmark is indicative of such use. For simple integer oriented benchmarks like those encountered in editing, compiling, and running operating systems, the YMP is ill suited since it is about the speed of Titan. This merely indicates that although a YMP is still faster for utility programs, it is not cost-effective, by almost two orders of magnitude and *should not be used as such*.

At first glance, the small Linpack case might look irrelevant to supercomputing. However, the average speed which Cray XMPs run at various large computer centers has in the past been equal to about 25 Megaflops/s, or the speed of Linpack 100X100 prior to Cray's recent compiler improvements. Note that for a single processor, it takes 12 times longer to get the same amount of work done on the distributed approach. However, the distributed approach is almost three times more cost effective or in principle, users spending the same amount could get three times as much computing done.

By automatically parallelizing Linpack even for the small case, the Cray YMP runs about 2.5 times faster using the 8 processors and has again become the world's fastest computer. Thus, we see the importance of parallelization to increase speed. Since the small Linpack benchmark is too small to run efficiently in parallel, the cost-effectiveness of the approach decreases over a factor of three. Since the Titan has only two, relatively slower processors, the effect of parallelization is not as great on cost-effectiveness. Stretch times in the order of 10-20 for the distributed, dedicated approach mean that even large users can get about the same amount of work done as with a centralized approach. Very large projects using a Cray center get only a few processor hours per day or about 1000 hours per year, while large users get an hour a day, and the average user gets an hour a week.

By using the peak speeds which are only obtained by running each of the processors at peak speed and in parallel, the difference in speed between the Cray YMP and the Titan is finally apparent. While the times stretch to almost 90 (i.e. to do an hour of computing on the YMP requires almost 90 hours on the Titan), the cost-effectiveness of the Titan still remains, but only by a factor of two.

Finally, using the graphic supercomputer, visualization is implicit in the system since each computer has a significant amount of power to render and display data. Modern supercomputing requires additional resources such as graphic supercomputers or high performance workstations just to handle the display of computed data. Future supercomputers must have embedded rendering hardware to provide both cost-effective and truly interactive graphics supercomputing since networks are unsuitable for giving users adequate interconnection bandwidth.

Politics, prestige, and the sociology of computing. With the entry of the Japanese into the supercomputer market, supercomputing has become an issue of national pride and symbol of technology leadership. While Japan now builds the fastest uniprocessors, the multiprocessor approach provides more throughput, and with parallelization more peak power. Table 1 indicates a very large number of companies have started to build high performance computers, including Steve Chen, formerly of Cray Research, and Burton Smith, formerly of Dennelcor. A recent report by the Office of Science and Technology

Policy urged the adoption of a government initiative for high performance computing, including a National Research Network (OSTP 1987).

When the first modern supercomputer, the Cray 1 appeared, university users and researchers used a highly distributed approach using the VAX 780. Several thousand VAXen were applied to scientific and engineering applications and only one Cray 1 was available at a university. The stretch time for the VAX was almost a factor of 100 over the Cray, while the cost-effectiveness of the approach was identical.

In 1983 with the emergence of the Cray XMP, the performance and cost-effectiveness of supercomputers was increased by a factor of 2 and 3 respectively. VAXen power was not increased, and only in 1984 did minisupercomputers with comparable performance/price to the Cray XMP emerge. By 1984 a strong national initiative based on the NSF Bardon and Curtis (1983) report and a plethora of supporting reports was started which resulted in the formation of the National Science Foundation's five National Supercomputing Centers under the Advanced Scientific Computing (ASC) program. The centers are now completely institutionalized into the NSF infrastructure at a budget level of approximately \$60 Million in FY89. The centers and the 5-10% of NSF researchers who use them have become a very effective lobby group for centers and for science.

State and federal governments including congress and large agencies delight in managing large scale, high-prestige national projects with opportunities for power and pork barrels. Similarly, large companies with large, central computing organizations enjoy the prestige of owning and operating supercomputers, independent of the payoff. At least one research laboratory believes that a supercomputer is essential for recruiting.

The purpose of this section is to present a more balanced view of the range of computing alternatives in light of the current emphasis by government and large organizations for centralizing computing.

Table 4. Central Cray YMP versus distributed Titan graphics supercomputer

	Cray YMP 832	Titan 24
Price	20.	.12
Processors	8	2
Mwords of memory	32	4
Dhrystones (integer-oriented)		
KDhrystones/s single processor	25	23
Dhrystones/s/\$ multiprogrammed	.005	.383
Whetstones (scalar floating point)		
MWhetstones/s single processor	35	6.5
Whetstones/s/\$ multiprogrammed	14	108
Linpack (100x100)		
Mflops/s single processor	79	6.5(12)*
flops/s/\$ multiprogrammed	31.6	108
Mflops/s parallel	195	9.4(21)
flops/s/\$ parallel	9.8	78

Peak performance (1000 x 1000 Linpack, theoretical peak)		
Mflops/s Linpack	2144	24(89) ???
Mflops/s/\$ Linpack	107	200
Mflops/s peak op rate	2667	32(83)
Mflops/s/\$ peak op rate	133	267
Millions of pixels rendered/sec	?	50

\*() the time stretch factor for Titan relative to the Cray YMP

### **Research Machines Which Are Entering The Market**

The following machines have come from DARPA's Strategic Computing Initiative, or single investigator research projects in basic computer science.

Systolic Processors. Kung's work at CMU on systolic arrays is beginning to pay off and arrays are being applied to a variety of signal and image processing tasks. The 10 cell WARP operates at an average of 50% peak for the problems of interest (speech, signal processing). This provides 50 Mflops for .35M (GE's price), or 142 flops/s/\$. Intel is building a single systolic processing chip, iWARP, that's capable of operating at a 24 megaflop rate. Such a chip would be an ideal component in a PC for vector processing in the '91 timeframe. Using the chip, a small board could compute at roughly 100 Mflops, and it would not be unreasonable to expect this to sell for 10K, or provide the user with 10,000 flops/s/\$. While the initial product was special purpose, the ability to use the WARP generally is improving with better understanding of the compilation process.

Text Searching Machines. Several companies have built specialized text and database machines. Hollaar has built and field tested a machine for very high speed, large database, text searches. The result to date is that enquiries are processed several hundred times faster than on existing mainframes, and the improvement increases with the size of the database since pipelined searching hardware is added with each disk. Teradata has a similar machine for databases.

Thinking Machines Corporation Connection Machine. The Connection Machine came from a research effort at MIT by Danny Hillis, which resulted in the establishment of the company to make the machine. Tens of machines are installed in a variety of applications from text searching, image processing, circuit and logic simulation to computational fluid dynamics. The current CM 2 model provides 64K processing elements, up to 1/2 Gbyte of primary memory, and operates at speeds up to 10 gigafloating point operations per second. Thus, the CM2 with 64K processing elements is the supercomputer for a number of applications. The Connection Machines are not multiprogrammed, but must operate on one problem at a time.

Multiflow and Cydrome as evolutions of the array processor. Based on Fisher's work at Yale on an extra wide instruction word to control a number of parallel (7 to 28) execution units, Multiflow Corp. was started up to exploit the compiler technology. In fact, the Multiflow can be looked at as either a SIMD computer with a small number of processing elements, or an extension of the traditional Array Processor, such as the Floating Point Systems computers. Multiflow's first product runs the Linpack benchmark at minisuper speeds and costs 1/2 as much. The relatively ill Cydrome

company built a similar product using ECL technology which provides even higher performance and better performance/price. One feature of this approach is that a compiler can exploit a substantial amount of parallelism automatically.

### **Experimental Machines In the Research Phase to Watch**

Berkeley and Stanford Multiprocessors. Both of these RISC-based, multiprocessor architectures are beginning to come into operation. So far, both have influenced commercial ventures both in risc and in multiprocessors. The Stanford project was the prototype for the MIPS Co.. chip design. The Berkeley chip designs were the precursor to SUN's SPARC chips. Another group at Berkeley has produced a first generation Prolog machine which has outperformed the fastest Japanese special Fifth Generation machines. The next generation Prolog computer is a multiprocessor/multicomputer to exploit both fine grain and message passing for parallel processing. Given the rapid increase in speed of POPs, it is unlikely that any specialization for a particular language will be able to keep up.

The University of Illinois Cedar Multiprocessor Project. Cedar is aimed at a multiprocessor with up to 32 processors in 4 clusters of 8 processors for executing Fortran in a transparent fashion and based on Alliant's FX-8. The prototype is will likely operate in 1989. Future work is aimed at more and faster processors.

The Very Large Multiprocessors: BBN's Monarch; Encore's Ultramax; and IBM's RP3. These three Strategic Computing Initiative (SCI) projects are all exploring the size and utility of large, multiprocessors, and provide over 1000 mips in sizes of 1000 @ 1, 128 @ 16, and 512 @ 2 mip performance, respectively. None have vector processing, and hence are unlikely to be used for mainline scientific and engineering applications requiring large numbers of floating point operations. However, the machines and automatic parallelizing compilers could provide sufficiently large amounts of power to attack new problems.

University of North Carolina's Pixel Planes is a scalable, highly parallel, SIMD architecture, providing the highest performance for a variety of graphics processing tasks such as solids rendering under varying and complex lighting situations.

ATT's Speech and Signal Processor. A large number of signal processing computer chips are arranged in a tree structured multicomputer configuration and provide over 250 gigaflops on 32-bit numbers. The machine fits in a rather small rack, and the resulting number of flops/sec/\$ is nearly one-million. The machine came from Columbia's tree structured multicomputer work and is part of the SCI.

The IBM GF11 and TF1 GF11 is a a SIMD computer with the goal of providing 11 gigaflops for the QCD calculation. TF1 has a goal of achieving 1.5 teraflop using 32,768 50 Megaflop computers

### **Computer Technology, Research and Training Agenda**

Packaging and Faster Circuitry. is nearly non-existent in university research. Very high speed processors require much better interconnection and packaging density.

Mass Storage. Given the extended performance machines, no radical improvements in size or speed are in progress to keep up with the processing developments described above that come from VLSI. A project which would couple say 1000, Gigabyte drives in a parallel fashion to provide an essentially random access of a terabyte of memory and utilizing a variety of specialized architectures would seem to be feasible. Such a system would be useful both as part of the memory hierarchy for the teraflop computer and as a database where high performance is demanded. Connection Machine's Data Vault, and Teradata's database computer are examples of what is possible in restructuring mass storage.

Visualization. In order to effectively couple to high performance computers, the scientific user community has recommended a significant research and development program in visualization (Visualization report 1987?). Today's supercomputers are capable of generating data at video rates and in order for humans to interpret data it appears the best way is to use direct coupled, high performance consoles. Two companies, Ardent and Stellar introduced graphic supercomputers based on this principle. Traditional workstation companies are increasing their computational abilities. While supercomputers and minisupercomputers currently rely on LAN connected workstations, it is more likely that both structures will evolve to have direct video coupling.

Room Area and Campus Area Networks (RAN/CAN). A RAN is needed to replace the various proprietary products and ad hoc schemes for interconnecting arrays of computers within the computer room and within systems involving high speed computation. At least three companies are building links and switches, using proprietary protocols, to operate in the gigabit range. The Hyperchannel is today's scheme, but the next generation must be a public standard. A combined FDDI and non-blocking, public standards based switch operating at 100 mbits/sec using fiber optics seems like a necessary first step which should last and evolve into the mid-90's.

Wide-Area Networks. Both intermediate speed (45 Mbit) and fiber optic (multi-gigabit/sec) switch development and research is non-existent. The networking dilemma is well defined (FCCSET 1988, Bell 1988). These networks are badly needed to interconnect the plethora of local area and campus area networks. Today, over 50 campuses have installed networks with an aggregate switching need of over 100 megabits per second, which implies an off-campus traffic need of 20 megabits per second to connect with the 1000-2000 academic, industrial and government research organizations. By making a system which could be used for both computers and communications, the two disciplines and industries could begin to become synergistic rather than antagonistic.

Addressing. The address limit of 32 bits on most of today's computers seems to be a severe constraint for every configuration except multiple computers. For example a solids data-set could easily have an array of 1000 x 1000 x 1000 elements, requiring a 33-bit address, assuming byte level addressing.

Dataflow As An Alternative Computational Model. Arvind's group at MIT has progressed to the point where it is building a dataflow computer that can potentially



outperform the largest supercomputer in problem domains with a high degree of parallelism and where vector computation techniques do not apply. Independent of the computer, a dataflow language may be the best for expressing parallelism in ordinary computers. Again, given the rate of increase in POPs, it is unlikely that an specialized architecture will be able to keep up with the mainline.

Neural Computing Networks. Various efforts aimed at highly parallel architectures to "simulate" the behavior of human processing structures such as the neuron continue to show interesting results.

Programming Is The Greatest Barrier to Use, And Requires Both Stability and Research. It is necessary to change the programming paradigm in order to get the highest performance from the myriad of new computer structures. However, the variety of programming models really isn't very large, given the variety of what would appear to be, different computers. The following table summarizes the main line of computational models and the corresponding near term computer structures which support the model:

Computation Model	Supporting computer structures
0. vector processing	x-supercomputers (one processor)
I. message-passing (coarse-med. grain)	workstation clusters (e.g. LAN'd SUNs) multi-computers (e.g. hypercubes)
II. multi/micro-tasking (fine grain)	shared-memory, multi-processors shared-memory, multi-processors
III. massive data-parallel	SIMD (e.g. Connection Machine)
iv. dataflow multiprocessors, and SIMDs	special dataflow computers,

Other computer structures such as the WARP (a pipelined array of systolic processors), neural networks, specialized SIMD computers, and the dataflow computer may ultimately require different computational models. In the long term, the above models may not be the best or even adequate to express parallelism. For now, we should build on what we know while simultaneously stimulating research on alternatives.

### **On Building The Teraflop Computer**

Two, relatively simple and sure paths exist for building a system that could deliver on the order of 1 teraflop by 1995. These are:

- a. 4K node multicomputer    800 Gflops peak
- b. Connection Machine    >teraflop with several million  
processing elements

Both machines require re-programming according to either the message passing or massively parallel data and single thread of control programming model. Given the requirements for users who want to exploit parallelism in the larger Crays will be starting down this path, it won't be unreasonable to assume that such a machine might be useful.

Today's secondary memories are hardly adequate for such machines. The lack of multiprogramming ability today may dictate using these machines on only one or a few very large jobs at the same time, and hence making the cost/job quite high, which would diminish the performance/price advantage. The cost of such machines will be comparable to the supercomputer of the 90's, or \$50-100M.

### **Will Applications Move Fast Enough to Fuel a Revolution?**

While it is difficult to predict how the vast increase in processing power will affect science and engineering generally, the following specific areas are clear.

Mechanical Engineering. Computers are being used in the design of mechanical structures ranging from automobiles to space craft and covering a range of activities from drafting to the analysis of designs including crash simulation for cars. Designer can also render high quality images and show the objects in motion with video. The vast increase in power should provide mechanical engineers with computing power to enable a revolution in mechanical design. Under this design paradigm every facet of product design including the factory to produce the product is possible without prototyping. We have observed this revolution in semiconductors and digital system in the last decade. Within the next decade mechanical engineering and companies could be transformed provided they respond. For starters it's better product quality. The big impact comes from drastically reduced product gestation times and the ability to have smaller organizations - which also translates to product elegance and quality.

Biochemistry, chemistry, and materials. With molecular modeling and computational chemistry, the design of molecules is now done interactively using large scale computers.

Large scale scientific experiments based on simulation. By having a dramatic increase in computational power, a range of system simulations involving many bodies would appear feasible starting with galaxies and going down to electron interaction to simulate the atom.

Animation. With the ability to compute realistic scenes, large scale computers provide an alternative to traditional techniques for film making.

Image Processing. Various disciplines including radiology rely on the interpretation of high resolution photographs and other signal sources. By utilizing high performance computers, the use of digital images and image processing is finally feasible. The use of satellite image data is transforming everything from military intelligence to urban geography.

Personal Computing. Today's large computers will continue to be used to explore the forefront of applications that will be feasible with the PC. For example, Ardent's graphic supercomputer, Titan, which sells for about one hundred thousand dollars is an excellent model of what will be available in 2001 at a price of less than \$6,000 (assuming continued price decline at the rate of 20% decrease per year). Many three dimensional phenomena from molecules to galaxies can be simulated at high enough speeds to transform modern science from experimental to computational simulation based. This

paradigm shift will transform every facet of science, engineering, and mathematics starting with education. Every home will have an almost unlimited laboratory to conduct experiments.

### **Some Suggestions for the Federal Government**

Diseconomy of scale continues to exist and favor small machines. Also, large regional computers can't be accessed effectively using today's limited networks, especially for interactive visualization. NSF's Advanced Scientific Computing (ASC) program is highly successful at supplying 5,000 researchers representing a few percent of the community with an average of 1-10 hours of supercomputing time per week with 1 hour the average. Smaller supers such as mini- or graphics-supers (VAX replacements) located in the labs would supply researchers an equivalent of at 10-40 hours per week and be under direct control of the research community. However most agencies have no way to support smaller, more cost-effective computers unless the prices are at workstation levels which can be supported by research grants. By not using small, dedicated computers under the direct control of the researchers more extensively, we are depriving researchers of a tool that should supply at least an order of magnitude more power than they now achieve through a shared super. One could imagine that this kind of infusion of processing power, directed at particular experiments could change the nature of science and engineering at least as much as the current ASC program.

The highly specialized computers offer the best performance/price and operate at supercomputer speeds, but cost more than a workstation. ATT's Speech processor that carries out 1/4 tera 32-bit floating-point operations per second is an excellent example that the gains possible by applications specific hardware and software. Again, agencies can't support them, nor does the community have the right combination of computer scientists and scientists and engineers working to tackle the problem of programming in any general way.

For the ultimate performance, SIMD machines such as the Connection Machine appear feasible. While the CM2 appears to be easy to program, *provided* the problem is rich in data parallelism, it doesn't support Fortran dusty decks that characterize scientific computing. NSF's research and computing Directorates don't support non-traditional supercomputing, given the large expenditure for the ASC program. Only now are the agencies with the greatest need (DOD, DOD, and NASA) beginning to examine the CM2 for computation.

If we want peak speed from any computer, programs will have to be rewritten to some degree to operate in parallel. One model using message passing where large amounts of data parallelism occurs, will work on nearly all high performance computers including: the Connection Machine, multicomputers, and multiprocessor supercomputers such as the Crays. The shared memory, multi-tasking model used by multiprocessors, including the Cray, doesn't work on the multicomputers since the computers don't share the same address space. The recent improvement in the Cray compiler to automatically parallelize as indicated in the Linpack benchmarks is a major technical achievement which will begin to allow general use and hence exploit the more rapid increase in available power through parallelism than through faster clocks.

The good news is that a vast array of new, highly parallel machines are becoming available. The bad is that only a few applications can be converted in any reasonable time frame to run on them, and furthermore no research programs or concerted efforts exists to solve the problem. This is simply not a case of new money, but rather a

redirection of resources.

## **A Challenge to the Computer and Computational Science Research Communities**

The computer systems research and engineering community is to be congratulated on the incredible computer performance gains of the last five years. Now is the time for the rest of the computer science community to become involved in using and understanding the plethora of computers that can be applied to the endless frontier of computational science.

The computer science has two options with respect to computational science:

1. continue to ignore the needs and challenge of scientific/engineering computing. This will certainly cause computer science to spring up in all the disciplines in order to deal with the interesting and hard questions the systems pose.

or

2. learn about the various forms of parallelism supported by the onslaught of new machines, by using and understanding them, writing texts, and training students (Current texts do not even deal with the "main line" of vector processing outlined above.) Understanding may enable work on automatic programming systems to analyze and rewrite programs for the above computational models. As a minimum, understanding will greatly facilitate exploiting the new machines.

### **Bibliography**

Althas, W. C, and Seitz C. L., Multicomputers: Message-Passing Concurrent Computers, Computer, August 1988, Vol. 21, No. 8, pp 9-19.

Bardon, M., and Curtis, K., A National Computing Environment for Academic Research. NSF Working Group on Computers for Research, National Science Foundation, July 1983.

Bell, C. G., "Multi's: A New Class of Multiprocessor Computers", Science, April 26 1985, Vol 228, pp 462-467.

Bell, C. G., A National Research Network, IEEE Spectrum, February 1988, Vol. 25 No. 2, pp 54-57.

Dongarra, J. J. Performance of Various Computers Using Standard Linear Equations Software in a Fortran Environment, Argonne National Laboratory Technical Memorandum, No. 23, August 20, 1988.

Office of Science and Technology Policy, A Research and Development Strategy for High Performance Computing, November 1987.

Misaki, E, and Lue K. M. Preliminary CPU Performance Study: The Ardent Titan and the Cray X-MP-1 on the Livermore Loops, Vector Register of the Institute for Supercomputing Research, Vol 2, No. 1, August 1988, pp 8-11.

McCormick B. H., DeFanti T A, and Brown M D eds, Visualization in Scientific Computing, Computer Graphics, Vol 21 Number 6, November 1987.

Smith, N. P. ed., Supercomputing Review, Issue No. 1, San Diego, Calif.

## **Figures**

Fig. 1. Simplified taxonomy of computer classes with representative computers in each class.

Fig. 2. Operation speed of various computer classes designed for high performance 1985-1994 (projected).

Fig. 3. Past, present, and future projection of the clock speed, number of processors and performance of Seymour Cray designed CDC and Cray Research computers.



## Extended Summary

Many, new high performance computers for scientific, engineering, and real time applications are emerging from DARPA's Strategic Computing Initiative (SCI) based on parallelism and gains in VLSI. These new technologies have stimulated both traditional companies and initiated venture capital backed start-ups to build the new computers. The research pipeline and transfer mechanism is clearly working for building high performance computers.

While power available from a single computer due to parallelism is finally increasing more rapidly than the leading edge hardware circuit and packaging speed gains (14% per year or roughly a doubling of performance every 5 years), this maximum power will not be available to a single job unless parallelism is conquered. Neither the user nor computer science community is moving rapidly enough to understand and exploit the potential performance gains (factors of 2 every two years) coming from an increasing number of parallel processors. Virtually all new computers provide, for the first time in the history of computer structures, a high degree of scaleability. With scaleability, a range of computer systems can be built from a common root processor, interconnection scheme, and common operating system which have basically the same architecture and programming paradigm.

The "main line" of scientific computing for the next generation (1988-1995) is likely to follow the supercomputer path based on a vector, multiprocessor with 4-6 and evolving to 64 processors. A range of supers (>\$10M), minisupers (<\$1M), and graphics supers (<\$100K) continue to evolve. No economy of scale, measured by processing operations/sec./\$, is observable over the range. The new class of graphics supers provides a dis-economy of scale for general purpose computing. The micro-super, a true personal super of Cray 1 speed in the \$10,000 price range will provide an even greater performance/cost improvement by the early 90's. The Cray 4, using 64, 2-gigaflops (billions of floating point operations per second) processors is targeted at 128 gigaflops in 1992. Today, reprogramming is necessary to achieve such peak power through parallel processing, although the Alliant minisuper and Ardent graphics super automatically parallelizes multiple vector-processors to work on a single task.

Plain Old one chip micro-Processors (POP's) using the reduced instruction set computer (risc) approach are becoming very fast and approach the speed for scalar/integer work of the largest mainframes and supers. Vector units can enhance uniprocessors to be more useful and cost-effective in workstations and small computers. A factor of 5-10 improvement in clock speed from the 20 Mhz to 100-200 Mhz range is feasible by the early 90's based on switching from CMOS technology to ECL technology. This evolution in clock speed is on the order of 70% improvement per year. By connecting these relatively, zero cost processors together in multiprocessors and multicomputers will virtually obsolete the style of computers supplied by today's mainframe and minicomputer suppliers.

Multiprocessors, sans vector processing, with a large number of micro-processors using the "multi" approach (Bell, 1985) have not yet proven themselves in scientific and engineering applications, due to lack of parallelizing compilers. However, such computers are clearly superior to existing "code museums" for time-sharing,

transaction processing, batch, and program development. Furthermore, a wide range of computers from two to several hundred processors can be constructed using the same basic components to achieve a form of scalability. Such computers providing several hundred to 1000 million, instructions per second will appear on the market in the 0.1 to 1.0 million dollar level.

Multicomputers, a collection of 32-1024 (or 4K) interconnected computers, each with its processor and memory, and communicating with one another via passing messages are the most cost-effective for single scientific jobs, provided the problem is compatible with the computer. Multicomputers require reprogramming, are used on one problem at a time (monoprogrammed), achieve supercomputer power, and cost \$25K-\$1M depending on the number of computers and their power. In contrast to general purpose computers, multicomputers are by program, made applications-specific.

A single instruction, massively large data (SIMD) computer, the Connection Machine - CM2 has become the supercomputer for a variety of applications including information retrieval. The CM2 is scalable and is currently available with 16K, 32K, and 64K physical processing elements, each with 8 Kbytes of memory. Like other applications-specific computers by program, the Connection Machine runs only one (or a few) programs at a given time, with a resulting performance/price advantage of a factor of at least 10 over a general purpose supercomputer.

A variety of other computer structures based on parallelism given in the simplified taxonomy of alternatives (Fig. 1) have proven themselves and are either emerging or show great promise. These include truly applications-specific hardware and software to carry out vision, speech, text, database, etc. tasks. By binding the application in hardware and software at design and construction time, a factor of 100-10,000 in some combination of performance or performance/price increase is possible.

A path to a computer capable of executing  $10^{12}$  floating point operations per second (flops) by 1995 looks possible utilizing either the multicomputer or SIMD (e.g. Connection Machine) approach. IBM Research has a project to build a multicomputer with 32K, 50 megaflop computers, the TF1.

Training and involvement by the computer science and engineering community to achieve high performance on real scientific and engineering applications is the main barrier to progress. The most formal training program is around the NSF supercomputer centers, but this effort is limited to simple vectorization of existing Fortran programs (i.e. dusty decks). Computer science has yet to embrace vectors as a machine primitive to be incorporated in texts and courses. An aggressive program to install, use, understand, write texts, and train students and researchers is needed beginning with embracing vector and parallel processing as a trivial, but necessary first step.

Long Range Plan for Scientific Computing ... has been submitted. Original plan is too ill-defined and open ended to manage. It is being segmented into smaller, more manageable and measurable parts. The role of the centers needs to be somewhat more than a provider of cycles, yet smaller than all the activities that could be addressed. It is difficult to express the notion of a central facility which is interconnected to a common facility and broad set of techniques (See attached.)

How does that vision fit with NSF overall view of academic computing needs? We are making a survey to determine the needs for scientific computing since the centers. This would include all facilities, workstations, networks, etc. We have an informal survey of computing today in academe for the scientific community and the FCCSET committee will take on this formally.

What areas will NSF emphasize in the CISE directorate? CISE includes a broad range of topics in computing, automation, robotics, and communications networking research.

In what ways has the computing community been involved in recommending priorities within the fields encompassed within Cise? Each of the 5 divisions have advisory committees. I have integrated the combined views.

How has the development of new more powerful and less expensive computers impacted the LRP for providing the academic community with advanced scientific computing capabilities? We have not changed the plan to reflect the fact that a new class of machines, the mini-supercomputer is available. We see it as being as pervasive as the minicomputer and an important computing for individual groups and departments. A very large fraction of all computing would be done on these machines, just as in the past. This would not diminish the need for centers. The only effect would be to increase the pressure for better interchange of programs.

How have developments in Hardware and software impacted the computational needs of the community? Unclear... continues to stimulate it to new heights.

Given changes in the computing field in the recent years, how does the Bardon Curtis report recommendation of 10 NSF s/c centers at academic institutions relate to your LRP? The need for the plethora of center is being satisfied in a number of ways: IBM, and the states.

Ultimate budget for sc in 87 was 5m below the original request, how was that reduction distributed and what were the consequences of the recommendation? Cut from von Neuman, Cornell, and hastened Phase I phase out.

How does that budget reduction compare to the levels of support estimated in the co-operative agreements with the National centers? Currently, the budget is below the amount we agreed to provide initially. We have been trying to find

ways to get more support at the centers in the form of industry use, computer industry support, state and university support. We are evaluating the support levels, various commitments, and the future needs.

What is the status of the phase i centers? They have been terminated as planned.

To what extent have the Japanese been successful in marketing their machines in the US? Only one machine exists in a university center at HARC. This is the SX2, the world's fastest computer.

What success have us manufacturers had in marketing their machine in Japan? There are no computers in Japanese labs or universities.

Are the Japanese unfairly cutting the prices of their machines? Don't know.

What is the administration doing in this area? Don't know.

What is the significance of separating networking from the ASC program?

What priority does NSF attach to networking?

What consideration has NSF given to the establishment of a national higher education computer network to support research and education?

## Types and Attributes of Centers

Center type	Examples	\$/yr	Duration test	Mission
STAR:Person/post docs Project/program	Mead Spur; CAD/CAM	>1M >2;>4	till un-productive till finished	eng.& science expts. breadboard, cabability
Uni.Facility (general) Uni.Facility/Discipline Nat.Facility/Discipline	Supers MRC's NCAR*, VLBA	5-10 2-8 10-50	forever obsolete; economics infinite	general facility for all facility and indiv. expts facil., researc staff and visiting expts scientists
Discipline-development Field/discipline suport	CER** ERC's	.5-1.5 1-3	discipline stable	develop a department people in a common area or descibed by common name

**Mission, sine qua non** (eg. build x, research field y, operate z)

### Structure attributes

"Centralization of resources" including: labs, fabrication staff, instruments, computers, programs, and databases (measure of theory vs experimentation)... not sine qua non (e.g. Berkeley Theory Center)

Additional Central Support and Overhead: fixed (including bricks and mortar, variable, non-productivity measures)

Control of research: local (via one time grants to center) vs. peer review as ticket to use the facility; includes industrial governance

Disciplinarity/coupling:

1. specific,
2. cross discipline,
3. general (all/any), and
4. multiple *completely independent* activities across one or more disciplines with no common output (e.g. Design, Communication, Semiconductor, and Systems ERC's)

Newness of area

Duration of centers: criteria for success and for termination; length

### Output attributes

Competitiveness, industrial coupling, and transfer mechanism(s)

Training, and curriculum development

Impact on advancement of science

\*National lab, which is "center" for university consortia is radically different of DOE National Labs due to various factors. Could these factors be the key to changing the utility of the DOE labs?

\*\*CER=shared use facility (hardware, maint. and support staff, and ovhd.)

Gordon Bell

11 February 1987

26 October 1987

Dr. Daniel E. Koshland, Jr. Editor  
Science  
1333 H Street, NW  
Washington, DC 20005

Dear Dr. Koshland:

Enclosed is a copy of the article or news story on a proposal for A National Research Network that I spoke to you on the phone about today. The title might be beefed up a bit as it is basically a proposal for a modern, superhighway system for information flow among computers and people in the 21st century. Senator Gore sponsored the original vision which the executive branch is well on its way to ignoring.

Given the Network's broad use by researchers in science and engineering, and crossing academe, industry, and government laboratories, I felt Science was a logical publication.

I hope you can look favorably on the story, as I think the scientific community badly needs the facility our FCCSET committee proposed. I'd be happy to live with a shortened constraint of an editorial if you think such a format would be appropriate.

Again, thank you for considering the article.

Sincerely,

Gordon Bell  
Assistant Director

Dr. Carl Ledbetter  
ETA Systems  
?, Minnesota

Dear Carl:

Paul Rotar just gave a report on ETA's progress, and it sounds like you are progressing toward being a viable supplier of scientific computers. It certainly aligns with my belief about how to succeed.

The most encouraging news is that you have a couple of UNIX ports nearing completion using just a few people. This shouldn't be a surprise since so many ports have been done over the last decade. Having Unix for Piper and the multiprocessors as the cornerstone of your operating systems is really important, especially since POSIX is becoming a government standard. For the good of ETA and its users, I hope the multi-headed operating system based on VSOS and Unix and providing backward compatibility doesn't get completed and marketed. I've never heard a really positive word about the operating system, except that it would presumably let users access their old files, but given the tiny base of 205's, backward compatibility can be solved in some way that doesn't penalize the future. A single system will save everyone at least 100 million dollars and years of grief. Go for one UNIX system and make it work!

I would hope the pipeline problem gets solved to improve the scalar speed beyond the 20% of a 205 as we know this is the number one rule (which CDC discovered) in building the "next" super-computer. It's always a mystery to me on how poorly, fundamental organizational knowledge (rules and myths) fail to get propagated to the actual engineers who unknowingly make the decisions. Neal has a book which apparently can't be published that should have this. (I would love to have a copy for my own use. Could you beg one from Neal?)

The nice thing about your processor is that it can now be used in a number of instances, and evolved with technology. Let me encourage you to really go all out in reducing the size (and clock period) so that it can be used in various future configurations. Spatially, I would think it could exist on a small card. Of course, having a computer which would sell for much less than a million dollars and used by scientists and engineers without massive bureaucracies and sales/support staffs is probably not possible within the ETA/CDC culture, but there are examples of companies who have made successful businesses building smaller machines. I still believe small computers help scientist and engineers substantially more than the large, big bang machines that really only operate best behind secure fences with only terminal access. (Cray already has that small market, and I see those users locked in and unwilling to retune their codes.) The problem, of course, with small machines is that the

government doesn't buy support them directly in the way we buy supers. Users have to really need and want them to buy them, and when they do, they get much more computing than from the pooled approach.

As for building fast machines, we simply have no evidence about using more than a dozen processors sharing a single memory in a single memory with automatic parallelization across loops (and this seems to require the right hardware for fine grain synchronization). For coarse grain, message passing parallelism, we have several examples of 1K computer systems (i.e. a processor and primary memory constitutes a computer) communicating through fast switches on a message passing basis. IBM's proposed TF1 supposedly has more like 32 K, >100 Mflop computers in a MIMD configuration. A smaller number of your machines would make it easier to reach a teraflop. This is one of the only two ways I see to achieving this goal... although it's unclear whether anyone can program it and achieve anywhere near the peak performance.

As a separate matter, I believe we need to be working on a standard, very fast, Big Room Area Network or Campus Area Network to interconnect supers to each other and to high speed workstations, to networks, and to back end database machines. I hope all the users have learned their lessons with the Network Systems Hyperchannel and will not go off and buy another proprietary interconnect. Is there a chance you folks would take the lead in this? How should it get developed?

Sincerely,

Gordon Bell

cc:

Erich Bloch

Robert Price

Paul Rotar



Mr. Seymour Cray  
Cray Research  
Chippewa Falls, Wisconsin

Dear Seymour:

Was glad to hear that the Cray 3 is coming along so well and that you are working on the 64 processor Cray 4. Also, it's gratifying to hear that you have a new head of software.

Although I understand (by trade press, rumour, and reported speeches) that Cray is only interested in building the most expensive computers, I believe it is possible to provide a line of compatible computers because of the basic technology today, without jepordizing the company focus. This has to be done carefully though, or you end up with another CDC. The approach CDC/ETA is using for building a one or two processor computer and their high end multiprocessor from a single module looks fine and really should provide a great deal to the vast number of users in small labs who have no access to large centers. I am convinced that having simple computers that any laboratory can buy and operate really helps science and engineering more than having the very large machines and bureaucracies that form around them. I don't believe in economy of scale, nor since the community is so slow in adopting to parallelism (i.e. reprogramming), there seems to be no way to exploit a 125 Gigaflops on a single problem. Similarly, the computer science community has been loath to work on the problem of helping the technical users.

Having seen the size and shape of the YMP processor modules, I would like to urge Cray Research to make either the modules or the design available to one of the companies building Cray-compatible computers in order to build a substantially smaller, one processor computer not requiring the very expensive installation and operation. A similar argument could probably be made based on using the Cray 3 modules. Alternatively, you might consider having a wholly owned division or company who use this approach but would not absorb the resources in a devisive fashion. It's conceivable that even a large company such as DEC would be interested in building small computers in this fashion.

As a seperate, but highly related matter, I now believe the teraflop computer that would work on a single problem can be built in the same time frame as the Cray 4. Two approaches look like they'll work:

- the large SIMD such as the Connection Machine It looks easy to program, given the problems have adequate data parallelism.

- the very large multicomputer formed by a collection of simple computers, interconnected by a high speed switching network for passing messages. We have considerable examples of several hundred computers operating this way now, but it's still unclear whether such an approach can be scaled up and built or whether enough applications fit the model.

I would be interested in learning more about the Cray 4 and how even 64 processors can be made to work together on a single problem... or if you are even going to try. The machine will still be useful as a package for 64 independent job streams, just as in the XMP. An alternatively, the message passing model can be used as in today's small hypercubes.

As a seperate matter, I believe we need to be working on a standard, very fast, Big Room Area Network or Campus Area Network to interconnect supers to each other and to high speed workstations, to networks, and to back end database machines. I hope all the users and manufacturers have learned their lessons with the Network Systems Hyperchannel and will not go off and buy another proprietary interconnect. Is there a chance you folks would take the lead in this? How should it get developed?

Again, forgive me for writing another letter for gratuitous advice on how to build supercomputers.

Sincerely,

Gordon Bell  
Assistant Director

CC:  
John Rollwagen, President

**Subject: The ETA P-series (Piper) Small Scale Supercomputer**

To: Director

CC: Assistant Directors , Division Directors (CISE), Al Thaler

From: Assistant Director, CISE

ETA (nee CDC division building the ETA 10) presented their new, small departmental supercomputer. The enclosed slides are relatively self explanatory. It is roughly a Cray 1 in speed, has a larger memory, and costs \$1 M and is easily installed. In contrast, a Cray may cost about \$1 M just to install (no machine) and it the power bill can run \$100K per year.

From my viewpoint, it is simply another very solid data point indicating a dis-economy of scale in supercomputing. On the next round of computers, I see this machine providing 30 floating-point operations per second per dollar, whereas the new Cray YMP and large ETA machines look like they'll only provide about 20. The small machine is only a factor of 2.5 slower, and hence with proper loading, should provide comparable or significantly better turn-around to a larger machine in a center.


The Advanced Scientific Computing program does not support this kind of computing at all, and again, I think science is losing. Our centers provide operations to the masses (about 5-10%) of the community, where the average user project only gets an hour a week of Cray time. A few users get from 500 to 2,000 hours. My own view is that the average user gets an insignificant amount of computing, and this won't result in a great change in the way science is done. A number of research groups need dedicated use - which equates to roughly 100 hours of Cray time or 2 orders of magnitude more computing time in order to really change the nature of science. This machine could easily provide the revolutionary change.

I know of no way to help the situation from our programs or from Washington. Unfortunately, I believe the insignificant amount of time we provide through the centers, to serve less than 10,000 users equates to at most 30-50 of these small machines.

I believe the research directorates should be looking at the situation on a discipline by discipline basis. A revolution is possible if NSF can find a way to support some of these machines at least at major centers and institutes from existing program and facilities budgets. Also, we should strongly encourage groups, departments, and universities to modernize their facilities with this new tool.

RECEIVED  
MARCH 24 1987  
MAIL ROOM

MEMORANDUM

TO: Gordon Bell  
FROM: Paul Rotar   
DATE: March 24, 1987  
SUBJECT: Center's Budget for FY88-FY90

---

The purpose of this memo is to inform you of my concerns over the level of funding provided for the centers in the FY88 budget request and to explain the funding needs in FY89 and FY90.

SUMMARY

The FY88 DASC budget request provides about \$45M for the Phase II centers Program. Under the assumption that CNSF support continues through the end of FY88 at level effort, the cooperative agreements project a need for \$48.5M. However, depending on financing options, \$54.5M to \$59.0M is needed.

Staffing levels at three of the centers are inadequate. In order that the centers operate at the service levels needed for Phase III, there needs to be an average increase of 10 FTEs per center for each of the next three fiscal years. This would provide an average staffing level of 85 FTEs. This staffing buildup in FY88 will add \$1.9M to the FY87 base with even larger increments in FY89 and FY90.

Equipment upgrades are needed at each center. Some have been planned and included in the cooperative agreements. Others are needed, but were not included. Upgrades in FY88 will add \$12.3M to \$17.8M to the FY87 base depending on the lease term. At a five-year lease term, the centers need \$54.5M for staff and equipment in FY88. At the optimum three-year term, \$59.0M is needed.

A comparison of the cost of operating each of the centers with NCAR's SCD shows that the average center support is about \$3M below that of a mature full service center.

Collaborative industrial programs have a long lead time and they have not been able to provide the additional needed funding.

Funds removed from the centers' program in the first years need to be restored so that each center has the resources to provide services at the Phase III and IV planning levels.

Cooperative Agreement and Actual Funding Levels

The following chart lists the FY87 and FY88 funding levels for the Phase II Centers Program. This chart represents only direct support of the centers and does not include the New Technologies program, any remnants of the Resources program, summer institutes, undergraduate education and the networking funds.

		CNSF	JVNC	NCSA	PSC	SDSC	Total
FY87	Coop Agmt	6.2	11.7	8.5	4.0	12.1	42.5
FY87	Actual	5.2	8.5	8.5	4.0	11.5	37.7
FY88	Coop Agmt	6.8	13.1	8.6	8.0	12.0	48.5
FY88	Budget						45.0
FY89	Coop Agmt	6.8	14.4	8.7	10.0	12.0	51.9
FY90	Coop Agmt	6.8	15.0	8.7	12.0	12.1	54.6

This table has the following assumptions:

- 1) Cooperative agreement amount extended at level effort through FY90.
- 2) No distribution of the funds among centers has been made at this time.
- 3) The CISE budget for DASC is \$48.2M, but \$1.6M is for New Technologies, \$1.0M is for education, \$0.5M is for summer institutes. \$0.1M is needed for costs related to program activities, leaving \$45.0M for the centers.
- 4) Assumes that cooperative agreement amounts are extended through the end of FY90; extrapolated at level effort.
- 5) CNSF-Cornell National Supercomputing Facility  
JVNC-The John von Neumann Center  
NCSA-The National Center for Supercomputing Applications (IL)  
PSC -The Pittsburgh Supercomputing Center  
SDSC-The San Diego Supercomputing Facility

Under the assumptions indicated, the FY88 CISE budget request shows a shortfall of \$3.5M from the cooperative agreement amounts.

An analysis of FY88 funding needs at the centers indicates that at a minimum, funding at the cooperative agreement levels is mandatory if these centers are to be robust, full service facilities able to operate on the leading edge. Lesser funding levels will prevent adequate staffing levels, CPU and hardware upgrades, and the establishment of collaborative research programs with industry. Even if funding were available at the cooperative agreement levels problems would exist. The cooperative agreements have underestimated the costs associated with full scale, leading edge, facilities because NSF's very aggressive negotiations with the proposed centers produced agreements providing funds below needed levels. There was expectation by NSF that high levels of industrial funding would offset any problems. However, with two years experience behind us and after much effort by the centers, it has proven very difficult to obtain significant contributions. I do not believe we can expect any rapid creases in industrial funding.

to quote from The Long-Range Plan for Advanced Scientific Computing, "Also it is important that the centers provide a unique service, not available at local academic computing facilities. We anticipate that many of the major research universities will soon be purchasing their own supercomputers. When that occurs, the centers will move into Phase IV, - the upgrade of their machines to the next generation (Class 7) of supercomputers." This Plan requires that advances occur at the centers in staff, hardware and excellence of service at the centers.

Staffing Levels

The activities in a generic scientific supercomputing center may be divided into a number of sections with the following staffing levels as follows:

Administrative & Secretarial	from 5	to 10
Systems	10	to 20
Operations	10	to 30
User Services	10	to 30
Advanced Methods	<u>5</u>	to <u>10</u>
Totals	40	100

The Administrative section is comprised of a director, deputy, administrative assistant, writer-editors and secretaries. The administrative assistant is a key person in budget management and financial details. In a large center, several writer-editors are needed and the secretarial staff is vital for the handling of telephones, travel, correspondence, meeting transcriptions, and conference logistics.

Systems staff maintain the operating systems on the main machines, front-ends and workstations. They also handle the systems responsibilities regarding the communications and networking functions within the centers. They perform any tailoring of the system to the local environment especially in the areas of accounting and communications. Two to three persons are required per major hardware item, i.e., the CRAY X-MP, VAX or IBM front-end, and common file system. An additional few are needed for communications devices and workstation software. In the latter area, someone with UNIX internals skill is needed. This definition of systems activity permits only maintenance. No wholesale development can occur.

Operations personnel not only operate the equipment, but frequently handle maintenance for small systems, communications gear, lines, etc. For twenty-four hour a day operation, seven days per week, it takes five FTEs to have one person on duty. Two or more are needed for safety considerations, i.e., no one is ever left alone for the sake of personal safety as well as equipment safety. Also when attending to the needs of nature the equipment is always tended. If there is a lot of data movement staff is needed to handle media.



Operations staff may also have to handle paper and microfilm/microfiche output. Some must be trained in film development and handling techniques. Others need skill in the maintenance of these devices. Workstations are now replacing film, but this will put stresses on the archival side of the activity because the results of graphics and visualization will be stored for future reference. Film was an excellent archival medium with easy storage (on a personal basis).

Maintenance of communications and terminal gear must be done in a timely fashion or there is instant dissatisfaction with the service.

User services staff perform consulting, training, circuit riding, documentation, and installation and maintenance of software libraries. The success level at the centers is very dependent on the skills of these people. User services personnel must have projects to work on that use the services of the centers. At a minimum, a half dozen consultants are needed. They need to be available more than eight hours per day since the centers serve a national constituency. Additional persons are needed for the other activities. During a software/hardware conversion a small staff can be overwhelmed.

The advanced methods group needs experts in computational mathematics, science, data bases and graphics. These persons form the basis for the advanced activities described below.

Centers staffs must grow with increases in the number of active users, with the complexity of the equipment configuration increases and as the user base extends to a larger geographical area. At the time of this writing (2Q FY87), the average staffing level at the DASC centers is 47. The centers now serve 3000 researchers in 200 Universities. Thus each person serves an average of 60 users. By the end of FY90, the user base should easily double. Hardware additions and upgrades are also planned. Together these will put pressures on the centers to increase the staffs to handle the workloads.

Current the staffing at all the centers include 40-45 or more persons that work for the centers in some capacity, but which belong to other organizations and either represent some cost sharing or contract work. This is not as effective as it might be if they were employed solely for the benefit of the centers and under their direct management. In short it effectively reduces the staffing available.

Staffing Levels

Staffing levels are expected to grow to between 80 and 95 FTES. Some fluctuation will occur depending on the budget and service requirements, but it is expected that center staffing levels should average 85 FTES by FY 1990. Following this staffing levels should more or less stabilize, declining slightly in years after major upgrades and rising somewhat in anticipation of upgrades. The following table projects my estimates of the necessary staffing levels: Let me reiterate that these levels are not affordable under current budgets.

	<u>FY87</u>	<u>FY88</u>	<u>FY89</u>	<u>FY90</u>
CNSF	43	56	70	87
JVNC	35	53	66	83
NCSA	74	80	85	90
IC	43	53	65	82
SC	<u>73</u>	<u>78</u>	<u>84</u>	<u>90</u>
FY Totals	268	320	370	432

An increase of 50 FTES is needed in FY88 over FY87. At an estimated FY87 average salary with benefits of \$50K/FTE and building in an inflation factor, this represents a need for \$2.7M plus \$1M for G&A expenses or a total of about \$3M. G & A now averages 43%.

Allowing for hiring delays, at least \$1.9M is required to get through FY88. The same pattern of increases and costs will continue through FY90.

The proposed staff buildup will permit the centers to provide a level of service and expertise that will distinguish them from centers where cycles only are available. Our Phase I Boeing center provided the fullest range of service and consulting and was the most popular center. Naturally it was also the most expensive Phase I center.

Our centers must offer services comparable to Boeing. Staff is needed to permit the development of leadership roles in generating new techniques and technologies for computational science and engineering.

While some of this can be provided through the normal grants program, kernels of expertise are required at the centers. Resident consultants and computational research assistants will provide the

nucleus for activity in the New Technologies, Computational Science and Engineering initiatives and other grants programs. Although there is not now a specific disciplinary focus at any center, a small group of researchers in one or two specific disciplines at each center would ensure that the software and services provided would enhance the scientific productivity for users of the centers and remove any impression that the centers are merely providers of cycles.

The centers should be the focus for the initiation of new algorithms and applications which will migrate to the university centers. They can serve as a driving force that stimulates all supercomputing activity in this country, but only if properly staffed. Current staff levels at CNSF, JVNC and PSC in FY87 will prevent any progress beyond their present activity.

#### Equipment Upgrades

All cooperative agreement budgets include equipment upgrade arrangements for all centers except SDSC. SDSC not only needs an upgrade in terms of computing capacity, but it also needs an SSD to enhance the interactive environment. The CTSS time sharing system operated only on disks on a machine without virtual memory spends a lot of time rolling job images in and out. In general CTSS overlaps this activity fairly well at SDSC because of restrictions on memory use, but this is improvable with an SSD. Also, there are I/O intensive problems that will not run well without an SSD. The SSD should be in the FY88 budget because it would help relieve the saturation at SDSC. SDSC also needs \$0.6M for mass storage upgrades in FY88. The CPU upgrade is needed by FY89. The SSD costs \$3M for a 128 Mwd unit, but it could be leased for \$1.0M per year including maintenance. In FY89, the CPU upgrade will cost \$20M or more. Allowing for a \$5.0M trade in and amortizing over five years, this adds \$3.8M to the budget beginning in FY89.

An upgrade at CNSF, from an IBM 3084 to an IBM 3090 400 with vector facility was provided by the vendor at no cost to NSF in October 1986. CNSF and IBM have worked out an agreement whereby in FY87, CNSF provides 12 systems and applications staff to support tailoring the software to Cornell's environment and providing feedback to IBM on NSF's software needs and IBM's software performance. IBM then provides a no cost upgrade to the existing IBM 3090. This would add 2 processors and increase both the main and extended memories to

the maximum available. Then in FY88 a second 3090 model 600 would be installed at an estimated cost of \$10M. A high-speed channel, 100MB/sec or faster, would link the two systems. It will provide the CNSF with about three times the current capacity and enable users to aggregate more processors on a single problem. This is an important upgrade because CNSF has the goal of making parallelism work. The \$10M could be financed over a period of five years at a cost of \$2.5M/yr.

The PSC budget doubles from \$4M in FY87 to \$8M in FY88 because cost sharing was front loaded and NSF furnished a CRAY-1S as a \$3M trade in for the X-MP/48. It increases another \$2M in FY89 to provide a funding level that will allow for an upgrade in FY89. The PSC budget assumed that the purchase price of the upgrade would be \$17M. At this time it is believed that this falls short by \$3M. To correct for this an additional \$0.8M per year is needed beginning in FY89.

The expected price increases for the upgrade also effect NCSA. They have level funding of \$4M per year for the Supercomputer mainframe. Assuming that they trade in the XMP/48 and refinance the balance along with the upgrade, they will need an additional \$1.5M a year beginning in FY89.

The FY88 upgrade at JVNC from an ETA-10/4 to an ETA-10/8 has a purchase price of \$8.8M. Amortized over five years, this has an impact of \$2.2M annually. Also JVNC needs funding for the ETA-10/4 which was removed in FY87. This comes to \$3M annually.

Upgrade needs for FY88 therefore total  $(0.6 + 1.0 + 2.5 + 4 + 2.2 + 3)$ . \$13.3M. With staff the overall increase from FY87 to FY88 should be \$16.8M. Thus the centers need \$54.5M  $(\$37.7M, (FY87 base) + \$16.8M)$  in FY88. The increase in funds from FY87 to FY88 of \$7.3M does not come closer to the amount needed to cover the additional staff and upgrades.

#### Optimum Financing Period

In this paper, a five-year payoff period was assumed for the mainframe upgrade discussion. However if the centers are to remain at the forefront of computing technology, payoff periods of three years for supercomputing equipment are more appropriate to permit the centers to obtain leading edge equipment. This is consistent with the January 12, 1987 statement to a subgroup of the PAC" the NSF Centers should exist with the largest peak power forever" The five-year amortization schedule understates the annual problem by \$4.2M. Also, the cooperative agreements would have to be extended for much longer periods to cover the financing.

### Comparative Costs

Another way of looking at the center's costs is to compare them with the cost of operating the NCAR Scientific Computing Division (SCD). In FY87 the NCAR budget has \$21.8M for science, \$13.8M for the SCD, \$9.6M for administrative activities. The cost to run SCD includes a proportional fraction of the administrative budget (\$3.7M) and for comparison with the DASC centers, it should include about \$0.5M for its share of the building. This gives a total of \$18M. As it now stands, the CISE FY88 budget request for the centers averages \$9.0M per center. Since NSF contributes about 60%, each center will now receive \$15.0M annually from all sources which is \$3.0M per center below the cost of operating a mature, full service center.

### Other Considerations

This memo concentrates on most of the large budget items. It should not be considered an exhaustive list. Improvements in disk capacity and performance may be desirable additions to the upgrades. Installation costs for upgrades can easily run to \$0.3M/site. As staff increases costs for space, furniture, travel, materials and supplies also increase. In FY86 and FY87 \$8.4M was removed from the program, i.e., funding fell below the cooperative agreements by this amount. If this continues as proposed through FY88 a total of \$11.9M as become unavailable to the centers. It is arguable that the money could not have been spent well in FY86 and FY87 because of startup problems in two centers. Now that startup problems have been overcome these funds need to be restored so that progress may be made. To continually fund the program at lower than needed annual levels will necessarily result in one of two outcomes. Either all the centers will be forced to lower desired staffing, service, and computing resource levels or NSF support must be withdrawn from one or more centers. There is little to be gained financially from the latter approach because NSF would have significant termination costs, including payoffs for existing machine loans, severance pay, building mortgages, etc. Also, we will lose leverage from all outside funding sources. Further there is no assurance that the monies saved after withdrawing support would be made available to the remaining supported centers. On the other hand, the inevitable result of a policy of continually funding below adequate levels is the ultimate withdrawal of NSF support from all the centers. At some future date this may

well be appropriate because the centers will have matured and become able to obtain sufficient non-NSF funding. However, at this time, this policy simply prevents them from meeting our mutual goals.

There is evidence that the failure to fund the program at the cooperative agreement levels has caused a distrust of NSF's willingness to commit to the future of the program. Certain entrepreneurial arrangements have become clouded and donations from now require matching funding from the NSF. The NSF is not the only organization that wants leverage and matching funds. There is less willingness for the Centers' other funding partners to provide more funds than before.

The collaborative industrial programs need years to become established and to grow. Only three substantive agreements are to be in place roughly two years into the Phase II programs: These are the NCSA-Kodak, the CNSF-Corning Glass and the SDSC-Aerojet General arrangements. The collaborative research route has long lead times. A policy prohibits the centers from selling time directly. This makes it difficult to quickly raise money to offset any NSF reductions.

I have heard expressions of dissatisfaction with the progress of the centers although our reviews last summer were favorable. It should be remembered that the Phase II centers have only been in operation, i.e., providing significant service since Jan 1986. This is only one year ago. The last two were on line toward mid 86. Only one (SDSC) was formed and managed by a trained cadre. The others were spun up using some help from the local Universities, but are largely under a management and staff that is learning. This approach allowed some dollar savings, but with some mis-step.

However considering their experience level, they have done well and are improving constantly. There are now significant scientific results from the program. Examples may be found in the reports the centers provided at the November PAC meeting. The PSC has published some of these in "Examples of Science Done at the Pittsburgh Supercomputing Center."

The Funding Summary FY88-FY90

The following table, based on a five-year amortization schedule, shows the centers' budget shortfall for FY88 through FY90. The expected budget level in FY89 and FY90 assume a 12.4% annual DASC increase based on FY87-FY88 experience.

	<u>FY88</u>	<u>FY89</u>	<u>FY90</u>	<u>Overall Shortfall</u>
FY87 base & inflation	38.8	40.0	41.2	
Staff	1.9	5.9	10.4	
SDSC Upgrade	0.0	3.8	3.8	
SDSC SSD	1	1	1	
SDSC MSD	0.6	0	0	
CNSF Upgrade	2.5	2.5	2.5	
PSC Upgrade	4	2	2	
PSC Upgrade Correction	0	0.8	0.8	
JVN Upgrade ETA10/4	3	3	3	
JVNC Upgrade ETA10/8.2	2.2	2.2	2.2	
NCSA Upgrade Correction	0	1.5	1.5	
Restore Transfer Fund	<u>0.5</u>	<u>0.5</u>	<u>0.5</u>	
Annual Requirement	54.5	63.2	68.9	
Budget Level	45	50.6	56.9	
Shortfall	9.5	12.6	12	34.1

A similar table based on a three-year amortization schedule:

	<u>FY88</u>	<u>FY89</u>	<u>FY90</u>	<u>Overall Shortfall</u>
FY87 base & inflation	38.8	40	41.2	
Staff	2.0	5.5	9.0	
SDSC Upgrade	0.0	5.0	5.0	
SDSC SSD	1.2	1.2	1.2	
SDSC MSD	0.6	0	0	
CNSF Upgrade	3.9	3.9	3.9	
PSC Upgrade	4	2	2	
PSC Upgrade Correction	0	1.2	1.2	
JVNC Upgrade ETA10/4	4.7	4.7	4.7	
JVNC Upgrade ETA10/8	3.4	3.4	3.4	
NCSA Upgrade Correction	0	4.5	4.5	
Restore Transfer Fund	<u>0.5</u>	<u>0.5</u>	<u>0.5</u>	
Annual Requirement	59	73.1	78.8	
Budget Level	45	50.6	56.9	
Shortfall	14	22.5	21.9	58.4



Conclusions

Funding is dependent on performance, but performance is also dependent on funding. Therefore, I believe it is imperative that the funds removed from the centers' program in the first years be put back in the later years beginning in FY88 to permit each center the resources to provide services at the Phase III and IV planning levels. The previous tables illustrate the benefits of such a restoration. At a five-year amortization, the restoration of \$11.7M covers the expected deficit. At a three-year amortization, it covers about half. Furthermore, such an increase would make the centers' annual operating budgets similar to those at older, well established institutions.



Subject: Alternatives to fund the centers up to their desires and abilities?  
To: Director  
From: Assistant Director, CISE  
Date: 21 August 1987

Based on the what I heard at a recent meeting reviewing the centers plans, it is clear that the ASC centers have the ability to absorb an arbitrary amount of funding for the following arguably reasonable things:

- incremental equipment for better balance (e.g. memory, disks, lines)
- upgrades to prevent center obsolescence
- increased capacity to meet industrial needs and opportunities
- "batch" and remote visualization equipment for movies

and lesser priority items including:

- courses
- grand challenges in computational science
- new technologies and new techniques in parallelism
- scientists to help in parallelization and visualization
- interactive visualization at the user level

We cannot possibly meet the requests. The disturbing fact is that **NSF is the sole source of support** at nearly all of the centers except Illinois, and the degree of support is increasing. Our current approach to funding has literally reduced industrial support. Except for Cornell and Illinois, the centers are really decoupled from industry; they are customers rather than research partners.

I would like to find some other ways to share this incredible funding burden. Here's my current list of options (ideas):

**1. Status quo.** NSF funds it all centrally, as we do now in competition with computer science. This is the worst of all possible worlds because the use of the facility is completely decoupled from the supply of the service. By being in CISE, nearly everyone associated with the budget, gets the **erroneous** conclusion that people working on computer science and engineering research have something to do with the centers. Little or no coupling or use of the centers is made by computer science. The machines aren't suitable for computing research, nor are adequate funds available for computational science.

If I make the decision to trade-off, it will **not** favor the centers, but rather centers will be funded at about the same as overall science.

- 2. Central facility.** NSF funds ASC as an NSF central facility. This allows the Director, who has the purview for all facilities and research to make the trade-offs across the foundation.
- 3. NSF Directorate use taxation.** NSF funds it via some combination of the directorates on a taxed basis. The overall budget is set by AD's. DASC would present the options, and administer the program.
- 4. Directorate-based centers.** The centers (all or in part) are "given" to the research directorates. NCAR provides an excellent model for say BBS, and MPS. Engineering might also operate a facility. I see great economy, increased quality, and effectiveness coming through specialization of programs, databases, and support. This is partially happening.
- 5. Co-pay.** In order to differentially charge for all the upgrades and incrementally nice facilities a tax would be levied on various allocation awards. Such a tax would be nominal (e.g. 5%) in order to deal with the infinite appetite for new hardware and software. This would allow other agencies who use the computer to also help pay.
- 6. Manufacturer support.** Somehow, I don't see this changing for a long time. A change would require knowing something about the power and throughput of the machines so that manufacturers could compete to provide lower costs.
- 7. Make the centers larger to share support costs.** Manufacturers or service providers could contract with the centers to "run" facilities. This would reduce our costs somewhat on a per machine basis.
- 8. Fewer physical centers.** While we could keep the number of centers constant, greater economy of scale would be created by locating machines in a central facility and running them more like LASL and LLNL where each run 8 Crays to share operators, mass storage and other forms of hardware and software support. With decent networks, multiple centers are even less important.
- 9. Simply have fewer centers.** but with perhaps increasing power.
- 10. Maintain centers at their current or constant core levels for some specified period.** Each center would be totally responsible for upgrades, etc. and their own ultimate fate.

**11. Free market mechanism.** Provide grant money for users to buy time. This might cost more because I'm sure we get free rides at places like Berkeley, Michigan, Texas and the increasing number of other institutions who do provide megaflops to their users.

I really question how we are going to fund this program in any fashion which permits the facility to be "traded-off" as part of a total research program. Only the disciplines can do this. I believe we should do the following:

1. consolidate equipment in fewer equipment-based centers to reduce cost and operate fewer physical centers at a greater economy of scale
2. have 3 or 4 directorate based centers and 2 or 3 general centers
3. use co-pay as a means to look at real need and as a way to fund specialized facilities such as 35mm movie equipment

Can I have your help on this matter?

## Why was it necessary to establish a central directorate for computing?

CISE recognizes the pervasiveness of the computer in society today and the unique opportunities for computing at this time.

The information society, which is <sup>about</sup> the largest sector of the economy, is based on computing and communication. Just as mechanisms were the basis of the industrial revolution, <sup>computer</sup> ~~computers~~ are the basis of the information revolution. *This*

Computing is found in virtually every scientific and engineering discipline as a base either as a tool or a component, and it is a science in its own right. In science, the Nobel Laureate, Ken Wilson, and head of the Cornell Theory Center housing one of the NSF supercomputer centers, explains computation as the third paradigm of science. The first being theory, and second, experimentation.

History has shown that government funding of computing research has been the main driving force of the revolution in computing, ~~that has become the largest industry today.~~

Finally, today, we have a ~~new opportunity~~ <sup>revolution - almost as important as the computer</sup> vis a vis parallelism to come off the technology evolutionary path of the last few decades that provide only x10 of performance per decade. *This rev. may be as important as the Computer*

## How Do you see Computing research affecting Competitiveness?

Directly through products. We have <sup>20A</sup> history of revolution that has occurred by <sup>ed</sup> funding university research - The Army funded Eniac and Edvac at Penn, the first computers that became the basis of modern computers, and the designers went on to create Univac. At MIT, DNR funded Whirlwind, from which came core memories, real time, air defense, air traffic control computers, interactive computing and the first computer aided manufacture. Digital Equipment Corporation came almost directly from the Whirlwind effort and team. Timesharing was first implemented at MIT; this became the basis of all modern computing. Graphics research, initially at the University of Utah, became the basis of all workstations and PC's, its how computers are beginning to be truly useful to everyone. ARPA funded communications networks for computing. The artificial intelligence-based expert system at Digital to specify how computers are put together was first prototyped at Carnegie Mellon U. This was the basis of the emerging AI industry. Universities are the main source of ideas and programs in VLSI design. *the only area we are leading in.*

Only this week, these example of NSF funded projects came across my desk: Don Knuth's program, TEX, is now the basis for modern typesetting of scientific and mathematics manuscripts. Two different parallel processing schemes, are now implemented by 5 companies. A research at

CER

- 22

- 2 more

- 1 for P<sup>2</sup>

→ 700k

→ phase out  
5 yrs  
recently

IUC's  
Industrial  
Univ  
Cent.

A study - Open in 83  
CS at Stanford  
13.5 M\$  
or 8.2 M\$  
for per year

*as a byproduct, it looks like it's been a new paradigm family program - literally program*



support and the contract was clear from the beginning that we were not going to continue support when our own phase II, centers became operational. The phase II centers are all now operational, ~~pretty much~~ according to our plan. By cutting our Phase II centers back 15%, would be a disaster; we simply can not maintain the systems at the performance levels we need, that is having the latest, and highest performance computer available on the market. This requires amortizing a computer over 3.5-5 years, the gestation time for a supercomputer. I do not support the concept of a "free market mechanism" for machine time at this time, whereby anyone can supply cycles. This mechanism didn't work and was the main drive why the government had to step in and form the ASC program in the first place. *at A 15% cut - costs could be maintained*

I am in the process of reviewing whether we have adequate funds to maintain our existing centers with the latest computers. It looks as if we are going ~~to have~~ to need more funds. I am not requesting more at this time, but believe we want significantly more help from computer suppliers, several of the states, some of the universities that host the computers, and industrial users. I believe the government is paying too much of the freight.

*We ~~also~~ have started to review the needs and compare this with a study conducted 3 years ago.*

### **Are you happy with the Program?**

Yes. We have 6K people on 2K projects, at 200 sites in all states. We see exciting results almost daily.

One of the great benefits from the program to date is the side effect of causing a number of great universities to acquire their own supercomputer. I don't believe any great university can afford not to have this kind of capability. For example, Berkeley has a small Cray XMP and an IBM 3090/200, Texas, Ohio State, Minnesota have or acquiring Crays. I hope the ETA computer will be successful, and replace the CDC 205's at various universities that have them. Michigan got the second 3090 after the center at Cornell. IBM has installed 40-3090/200's which could supply significant computing power (each processor of a dual is about equal to a Cray 1). In fact today, I estimate that we have the equivalent of over 110 times the Cray 1 available to the university research community, about 70 of this is in unis, about 40 at the centers (including NCAR) in 4 Cray xmps, 1-Cray 1, 1-3030/400, and 2-CDC 205's that are to be replaced with a machine of 20x a Cray 1.

As an alternative, three companies are building and installing mini-supers, all of which can do many of the tasks supers can do on a cost-effective basis. Many more designs are in the wings.

As the person responsible for getting about 3000 computers of the minicomputer price class into the scientific community in the form of

VAX, I think the future will give us lots of options in the way to do computing. Today, supercomputer users generally access supers at the end of a very slow network. This limits their own abilities in a different way, particularly in being able to visualize results. Many things (other centers, superminis, and networks) have changed since the establishment of the centers program, and we must continually evaluate the options for the future.

In all scenarios, I continue to see the need for a few centers which have the latest and fastest computers.

Utah has just implemented a text searching scheme that promises to be able to retrieve any text in any size database in virtually 0 time. Kamakahar's algorithm at BTL came out of extension of his thesis work at UC/B.

The supercomputer centers produce results regularly: America's cup, Kodak Material, Corning (1/6 throw away) simulation of the new superconducting materials, search for cold virus serum, molecular modeling and computational chemistry, we even have work to use the computer as a computational telescope.

Computers are critical to CAD, CAM, CAI, CAI, ... in every environment from home, office, laboratory, vehicle, or factory. We especially are focusing on hardware in this budget, note the increase in the MIPS area.

Finally, we still have a + balance of trade in computers, but its fading fast. Japan is breathing down our neck in every phase of R...D, and every area from AI to payroll.

Bottom line: We have no trouble in measuring results, including gestation times. It is quite rapid, and it can and must be even faster.

### **What are you doing to help the education process?**

I ~~mainly~~ believe that the big force that drives the education process comes from the right balance between research and teaching by first rate researchers. Much of research comes from student questions.

Let me give a homely example of the interaction of teaching and research. I took 6 years off from Digital to teach and do research at Carnegie Mellon university from 66-72. I wanted to explain how simple computers were and to have computers design them. We came up with 2 notations, that later became languages to describe, simulate and ultimately now to begin to automatically design computers. The text we wrote is still a classic on computers, and many simulators use the language, and at least one company sells the program. All of this came out of a research direction and drive that was largely pedagogical.

I also believe that the work we are doing indirectly in CAI in some of the leading universities will ultimately filter into all forms of education.

### **Are you familiar with Rep Sabo's Proposal to have NSF fund the Phase I centers by cutting 15% from the Phase II budgets, and then ultimately go to a free market for all supercomputer service?**

The Phase I program was established to buy computer time from various organizations, including three companies and three universities (Colorado state, minn, Purdue) who had supers. We had no long term commitment for

Long Range Plan for Scientific Computing ... has been submitted. Original plan is too ill-defined and open ended to manage. It is being segmented into smaller, more manageable and measurable parts. The role of the centers needs to be somewhat more than a provider of cycles, yet smaller than all the activities that could be addressed. It is difficult to express the notion of a central facility which is interconnected to a common facility and broad set of techniques (See attached.)

How does that vision fit with NSF overall view of academic computing needs? We are making a survey to determine the needs for scientific computing since the centers. This would include all facilities, workstations, networks, etc. We have an informal survey of computing today in academe for the scientific community and the FCCSET committee will take on this formally.

What areas will NSF emphasize in the CISE directorate? CISE includes a broad range of topics in computing, automation, robotics, and communications networking research.

In what ways has the computing community been involved in recommending priorities within the fields encompassed within Cise? Each of the 5 divisions have advisory committees. I have integrated the combined views.

How has the development of new more powerful and less expensive computers impacted the Lrp for providing the academic community with advanced scientific computing capabilities? We have not changed the plan to reflect the fact that a new class of machines, the mini-supercomputer is available. We see it as being as pervasive as the minicomputer and an important computing for individual groups and departments. A very large fraction of all computing would be done on these machines, just as in the past. This would not diminish the need for centers. The only effect would be to increase the pressure for better interchange of programs.

How have developments in Hardware and software impacted the computational needs of the community? Unclear... continues to stimulate it to new heights.

Given changes in the computing field in the recent years, how does the Bardon Curtis report recommendation of 10 NSF s/c centers at academic institutions relate to your Lrp? The need for the plethora of center is being satisfied in a number of ways: IBM, and the states.

Ultimate budget for sc in 87 was 5m below the original request, how was that reduction distributed and what were the consequences of the recommendation? Cut from von Neuman, Cornell, and hastened Phase I phase out.



How does that budget reduction compare to the levels of support estimated in the co-operative agreements with the National centers? Currently, the budget is below the amount we agreed to provide initially. We have been trying to find ways to get more support at the centers in the form of industry use, computer industry support, state and university support. We are evaluating the support levels, various commitments, and the future needs.

What is the status of the phase i centers? They have been terminated as planned.

To what extent have the Japanese been successful in marketing their machines in the US? Only one machine exists in a university center at HARC. This is the SX2, the world's fastest computer.

What success have us manufacturers had in marketing their machine in Japan? There are no computers in Japanese labs or universities.

Are the Japanese unfairly cutting the prices of their machines? Don't know.

What is the administration doing in this area? Don't know.

What is the significance of seperating networking form the ASC program? *... little more*

What priority does NSF attach to networking? *... priority ...*

What consideration has nsf given to the establishemnt of a national higher education computer network to support research and education? *...*

*It's raising the priority ...*

Comparison of Super/Mini-super

	Supercomputer		Mini-supers
μM/yr. - NSF	11	5yr. amort.'88	0.25 .15/yr@5yr+.1 oper
People served	1200		25 5-300 users
Power (Cray 1s)	12	84	0.5 3/86
\$M/power	0.92		0.5
Stretch time	0.33 for 1 processor		2 vs. Cray 1/S
Power (t+1)	36 3/88; YMP		1.5 1/89 wild guesstimate
\$M/pwr(t+1)	0.31 too low		0.17
Network \$(M)	1.2 at 1,000/user		0 Net costs not incl.
Discounts	0 Norm.		0 20% typical
Staff size	60		2 0.1 operational costs
	Case 1		
Constant power	12		12 for 1200 users
Machines	1		24
Cost per year	11		6
Users/machine	1200		50
	Case 2		
Constant users	1200		1200 25 users/machine
Machines	1		48 *5 centers=240 Sites!
Cost per year	11		12
Pwr. hr/yr/usr	72		144
	Case 3		
Constant budget	11		11 for 1200 users
Machines	1		44
Total power	12		22
Pwr. hr/yr/usr	72		132
	Case 4		
Constant budget	11		11
Users	1200		1100 25 users/machine
Pwr. hr/yr/usr	72		144
Advantages:	Large staff	Users involved	more people learn
	Supports the	Grand Challenges	Science-based management. Need based
	Nec. for 10%	of users	Serves most of the user population
			and physical memory
			with W/S access
Disadvantages	Large staff	;	Prog support with users... they learn
	Central NSF	staff	Can't support Grand Challenges
	Small Virtual	& physical mem.	
	Lack of UNIX	& W/S comp. O/S	
	Network access	\$ & poor graphix	

Full Center Staff; 100K/yr mini operations staff

## MEETING THE GRAND CHALLENGES OF SCIENCE AT THE NATIONAL SUPERCOMPUTER CENTERS

### IN THE BEGINNING . . .

In the two years since the National Science Foundation announced the establishment of National Supercomputer Centers, much has been accomplished.

The National Centers have acquired, installed, and are operating their supercomputer systems and have obtained staff for management, operation, networking, provision of user services, and new projects such as visualization. The Centers are fully operative -- some centers are already saturated and total usage has been growing rapidly. In the brief two year period, researchers in all scientific and engineering disciplines have gained access to a wide range of advanced computing capabilities. At this time there are more than 3000 scientists using 3 CRAY X-MP's, a Cyber 205, and an IBM 3090 at the National Centers. We serve virtually every discipline of science and engineering.

The Centers have been successful at attracting non-NSF support to help in these efforts. This includes (i) donations of computer systems and workstations from the computer industry, (ii) corporate support for industry/university collaborative research, (iii) state and university support for networking and communications, workspace for staff, renovations, and new structures, and (iv) industry discounts on equipment and software. The NSF Centers have many means of access. Several of the Centers have established consortia networks to support remote access. Some Centers have used more ad hoc methods - "800" numbers, packet-switching networks, etc. The Centers also continue to support rapid development of NSFNET and have assisted in its early stages of implementation.

### AND NOW TO THE FUTURE . . .

The best and proper role of the Centers in the future is the key issue in establishing plans and budgets. In December 1986 the Directors of the six NSF Centers (including NCAR) met to discuss future directions and the resulting program needs. The Directors agreed upon the following goal for the National Centers:

"To advance science and engineering as far as the cutting edge of computational technology will allow."

The purposes of this paper are to express the means by which the Centers will address this goal and to seek from the Foundation the level of funding required to make it possible.

#### TO ADVANCE SCIENCE AND ENGINEERING . . .

The Centers have created an intellectual and computational environment that encourages bold attacks on truly important problems. They are now attracting researchers in universities and industry to collaborate on interdisciplinary problems of national scope and with common computational needs. To fully realize the potential of the Centers, there must be substantial, long-term support at the national level for a broad intellectual and computational environment.

We see emerging several examples of very significant national and interdisciplinary problems (which we choose to call Grand Challenges). Below we discuss, in general terms, the requirements for achieving major successes in these areas. We follow this with an assessment of what is needed in order to help the centers meet these requirements. The problems are:

1. **Electronic Structure** - the binding together of atoms to form molecules and material substances underlies many disciplines of science and engineering (chemistry, biology, materials science, etc.) and many industrial areas. The Schrodinger equation accurately describes this binding. In practice, the cost of simulation rises rapidly with the number of electrons. Supercomputers are essential for greater than 4 to 100 electrons (depending on the algorithm and degree of reliability one seeks). A rich set of algorithms exists for the study of electronic structure. Nevertheless, major technological and algorithmic advances and novel computational strategies are needed to go much beyond this level. Advances here can have a major impact on the technological base of our society.

2. **Molecular Biology** - accessibility of supercomputers has enabled scientists working in the area of macromolecular modeling to make significant progress toward a fundamental understanding of the nature of the molecules of life. New breakthroughs have already been made in the structure and analysis of a human cold virus and in quantitatively predicting the relative enzyme-substrate binding energies using an empirical potential energy model. These advances hold great promise in the area of pharmaceutical chemistry. Supercomputer access has also rapidly accelerated directed efforts toward rationalizing three dimensional structures of proteins on the basis of their amino acid sequence. This area of research, known as the "protein folding problem" is fundamental to an understanding of structure-activity relationships in biological systems. The more powerful supercomputers of the future will still be challenged by the size and complexity of such studies.

3. **Global Geosciences** - observations over the last few years have revealed striking changes in key indicators of the state of the planet: the atmospheric

concentrations of carbon dioxide, methane, fluorocarbons, and ozone are all changing dramatically, indicating that the dynamic balances that control the planetary environment are being altered. We see an emerging image of profound change occurring on the scale of decades to centuries. In contemplating the immediacy and threat of change on the global scale, science has discovered anew that processes on vastly different spatial and temporal scales are all interrelated in the components of the Earth System--the atmosphere, ocean, land surface, and the terrestrial and marine biospheres. All respond to variations in the solar energy that drives the system, and all are part of the complex process of controlling the Earth's energy budget and the planetary environment.

The challenges of understanding and predicting global change are not only straining the capabilities of the most advanced supercomputer systems, they are forcing re-examination of the traditional roles and structures of the Earth Sciences, of the aims of individual scientists, of the organization of university departments. The scientific community, federal agencies, and governments throughout the world are responding with initiatives such as the World Climate Research Program, the proposed International Geosphere Biosphere Program: A Study of Global Change, the NASA Earth Systems Science Program, and the NSF Global Geosciences program.

4. Computational Fluid Dynamics - The nonlinear, sometimes chaotic, aspects of flow processes are pervasive in science and engineering. Much progress has already taken place, including the development of the theory of solitons and other nonlinear waves and the analysis of routes to chaos in dynamical systems. Computational studies have been particularly effective in analyzing these problems and extracting new ideas. By supercomputer simulations scientists and engineers can, for example, perfect the design of vehicles to reduce drag and increase flow mixing to improve combustion and heat transfer.

Similar advances may be expected in such wide ranging problems as methods to predict oil reservoir dynamics, ocean surface wave interaction, atmospheric data simulation, and environmental engineering. Even modest improvements in these problems are worth billions in potential cost savings each year. Problems of major interest are still waiting, however, for more powerful supercomputers to make them tractable - to simulate full-scale aerodynamics for the next generation of aircraft, for example, the computer system must be capable of producing one billion operations per second on a sustained, real performance basis

5. Decoding Genomes - Advances in laboratory techniques in mapping and sequencing genomes of living organisms are challenging the ability of computational resources to keep up with rapidly increasing needs in genetic modeling and the coming flood of detailed genetic information. Within two years, Japanese devices are projected to be capable of sequencing over 1 million base pairs per day. The community is discussing the technical possibility of sequencing the entire human genome (3 billion base pairs) within the next decade. Some believe that it would be better to concentrate on lower



organisms first. Regardless of the outcome, there will be an explosive growth in the years ahead of sequencing of genomes.

The very nature of this research is inherently interdisciplinary, drawing on many subfields of biology, as well as numerous aspects of computer science and computational technology. The data storage and computational requirements will strain today's supercomputer systems. Clever new algorithms and computer architectures are needed. Clearly a national network is needed to tie together the laboratories producing the new data, the biologists using the data, and the national centers (such as the Los Alamos DNA database) storing and processing the data.

6. Brain Mapping - Sir Francis H. C. Crick observed, "There is no scientific study more vital to man than the study of his own brain. Our entire view of the universe depends on it." Basic information is still lacking in such areas as measuring and displaying individual differences among human brains; organizational changes as affected by genetic factors, human growth and development, and trauma; dynamic processes - electrophysiological, metabolic and biochemical; and structure-function relationships as they relate to perception, judgement and behavior. It will be a major challenge to develop a system capable of mapping in three dimensions and displaying the entire human brain at microscopic levels of detail. The demands for storage, manipulation and processing vast amounts of information are immense but overshadowed by the need to develop algorithms to identify and quantify brain structures, modeling of dynamic processes, and developing relationships helpful in diagnosis and treatment of disorders.

7. Modeling the Global Economy - As our planet moves toward an integrated global economy, the need for realistic computational models of this complex system becomes more acute. There has been progress in particular models (Project LINK, Input-Output, FUGI, etc.) which take a certain mathematical framework and interact with data bases. However, each model captures only a part of the interactive system. What is needed is a major interdisciplinary effort to create a super-model which couples together the key processes which we know to be important, such as demographics, resource distribution, monetary flows, transportation and communications networks, trade flows, variations in political and economic systems, etc.. This will require bringing together for extended periods leaders in many different sub-fields of economics and political science, with experts in mathematics, algorithms, and the science of complex interactive systems. In addition to reforming the mathematical model, vast holes in the databases representing the real world need to be filled. Such a goal obviously will require international cooperation and a decade or more of work.

#### THE ROLE OF THE NATIONAL CENTERS IN THE GRAND CHALLENGES

We find that there are some common prerequisites for achieving full benefits of supercomputers in these Grand Challenges. First, each of these needs a fully interacting computational community from a wide range of disciplines to ensure that the best algorithmic and scientific approaches are identified and pursued. Second, each needs continually evolving facilities at the cutting edge

in capabilities and performance to ensure the best use of algorithms and timely solution of problems.

The key roles we see for the centers in addressing the Grand Challenges are: First, they provide a concentration of technical and computational expertise capable of supporting these projects. Second, they provide major resources at the leading edge of computational technology, that have to be centralized for maximum capability and performance because their cost is and will remain beyond the reach of most universities. Third, they provide a flow of information between various computational communities solving these problems - avoiding duplication of effort and speeding the widespread use of new information and computational techniques. Fourth, they interact with industry through collaborative research, training and the communications inherent in the U. S. scientific community, in productive ways that enhance technology transfer.

There are many applications of supercomputers that require the full range of human and technological resources of a national center. The centers will have a critical role in inaugurating new scientific and engineering applications and will transfer these technologies to the industrial and university computational science and engineering communities. This will contribute to strengthened United States competitiveness in the world marketplace.

#### ... AS FAR AS COMPUTATIONAL TECHNOLOGY WILL ALLOW

The National Centers are ideal for commissioning tomorrow's supercomputers. Once these systems are acquired, the Centers will operate them, establish a user base, and install vendor developed enhancements. Scientists will benefit by early access to systems of the future and the computer industry will benefit by favorable exposure and user input on projects in development.

There are desirable features of computing which are of great importance to the computational science community but which have not yet materialized. These include higher bandwidth I/O channels, interactive graphics, improved mass storage systems, and standards. The Centers are, individually and collectively, attempting to influence the computer industry on such future developments.

The economic competitiveness of the United States requires rapid advancements in our technology and this, in turn, requires a steadily new and increasing supply of technical experts. The Centers are one of the very few resources available for training and educating students and young investigators in supercomputing and advanced computer technology.

WE CONCLUDE ...

The five National Supercomputer Centers have come into existence in a very short time. They all have achieved major leverage from funding outside the NSF - \$150 million estimated over five years out of a total funding of \$370 million. Still missing to support the above Grand Challenges and the opportunity to push technological boundaries are:

1. Adequate human resources to pull the scientific and engineering computational communities into full usage of the supercomputing environment, and to train a large number of researchers in the expert use of computational techniques,

2. Achievement of balanced computational resources at the centers(balanced among cycles, memory, graphics, networking, mass storage, etc),

3. The financial strength to maintain this balance at the leading edge of advanced technology, as required by the Grand Challenges, and

4. Far reaching, high speed networks to provide researchers effective access to these resources.

Because of the immense national importance of computational science and engineering to the competitiveness of the U. S. in the international marketplace, we request that the NSF increase support for the National Centers. This will improve the computational facilities and staff at the Centers, enrich the intellectual environment, and facilitate high quality access to these facilities. Thus the National Centers will be able to fulfill their role in advancing science and engineering as far as the cutting edge of computational technology will allow.





OFFICE OF THE  
ASSISTANT DIRECTOR  
FOR COMPUTER AND INFORMATION  
SCIENCE AND ENGINEERING

## NATIONAL SCIENCE FOUNDATION

WASHINGTON, D.C. 20550

January 26, 1986

Dr. Jack Worlton  
3089 Villa  
Los Alamos, NM 87544

Dear Jack:

I enjoyed your very thoughtful and useful comments on my paper, your slides, and your paper on Computational Mechanics. I urge you to submit the paper to either the ACM Communications or IEEE Computer Magazine in order to get the computer science/engineering community involved in computational science. Let me also encourage personal interaction with this community with lectures and papers.

The lack of communication between the traditional users of supercomputers and the computer science community which trains new users and computer systems builders, and studies "realistic" next generations of machines, is a most difficult problem for NSF. I believe the communications gap between the two communities is impeding progress on all fronts from applications software to basic systems software. It will certainly limit our ability to use parallel computers for scientific problems. As you may recall in the three unsuccessful attempts at LLNL, LBL, and LASL, it is non-trivial to build parallel computers. Yet the mainline of computers is clearly the shared-memory, vector, multiprocessor for all classes (supers, mainframes, minis, and workstations). I don't believe the answer is to ignore the computer science community as users attempt to build parallel machines, but rather to involve them. NSF's initiative in Computational Science and Engineering is aimed at doing this.

Your comments on the paper stimulated these thoughts and questions.

### ABOUT THE DATA AND THE DIAGRAM

**100 x 100 Linpack (or any other single number) Benchmark.** If I am allowed only one number to characterize scientific computer performance, it would be the 100x100 Linpack benchmark. I have found good correlation with two other benchmarks, and that it is the rate that the Cray's I've asked about run, averaged over long periods, e.g., days. I should have put the ranges on performance, as you did on the plots, but I would sure like to understand how the distributions of speeds for various applications vary by environment and user sophistication (e.g., National Labs, NCAR, NSF Centers, and traditional computation centers). The table does have the 300x300 performance metric, and the derived performance/price metric. The reader should draw conclusions from the range numbers.

Note, a recent report by a committee sponsored by Jim Decker's FCCSET committee argued for better understanding of benchmarks. The report suggested that DOE and NSF carry out this work. Surely the evaluation techniques at LASL can shed more light on performance as it has over 40 years experience in scientific computing. I would like to get our community involved in working with the National Labs on this as soon as possible. Understanding virtual memory and cache performance is also critical as IBM comes into the market. I believe we'll find out these techniques actually work, despite the fact that Cray doesn't use them.

**Configuration Size (i.e., cost).** The secondary memory is expressed in Gigabytes. The IBM configuration has 60 Gigabytes, which implies the Cray is artificially lower priced, especially since many of them also have IBM disk farms as back-ends, and they all have mini front-ends. It is also unclear what the appropriately scaled disks for the mini-supers and workstations should be, but I attempted to put down what I thought were operational configurations, and that's why I have several configurations for the SUNs. Having appropriate secondary memory and the full i/o complement for all the supers and alternatives in the other machines clearly turns the plot into the areas as you suggest.

**Time discounted performance metrics.** In plotting computer price and performance evolution, the key variable is time. The plot doesn't give these, but the table does give the delivery dates so the reader can make his own adjustments. A sophisticated reader would discount machines at about 14% per year (a factor of two in 5 years), the technology evolution factor for current supercomputer circuit/packaging technologies. The Cray 48 (shown erroneously in its recently faster clocked version with the initial ship date) would be more impressive if I were to use "discounted" metrics in the comparisons. Not discounting price or performance with time probably creates the biggest error in computing performance/price.

**The bottom line** is that for the table, the reader has the caveats (especially about ETA). It is from the table that I draw a few conclusions, for example, the need for standards, the need for networking to make the centers most effective, the plethora of alternative computers creating a hierarchy and acting as "feeders" to the centers and, of course, the need for more understanding about use, etc.

I am not unhappy with the numbers or the derived numbers except that they may underestimate IBM. If I update or submit the paper for publication, I would like to use your plot. The reader has to be able to interpret the table, ranges, and understand the great variations due to systems management, use, networking, and configuration.

## ABOUT THE CONCLUSIONS

I agree, supercomputers have the highest potential performance, and if people exploit them with proper vectorization, then the performance/cost can be much better than the uni-processor mini-supers. Certainly the variation is higher because you multiply the variation due to vectorization of a single processor by the number of processors. This point can be understood in looking at the 4 processor XMP which runs, on the average, at Linpack 100x100 speed (4 x 27). Yet the hand-coded peak Linpack (713) can almost reach the peak of the machines. Do you see wide variations in the use of your machines at LASL depending on user sophistication, degree of time-sharing, type of problem? Again, the issue is that the best cost-effectiveness of the super comes from using it at its peak. Given the performance needs at LASL and LLNL, are you not using smaller machines in the super region? Are you using some of your XMP's in a parallel processing fashion (either with shared memory or with message passing) to get into their peak range?

Could it be that on the average, the supers should be used with front-ends for non-vectorized work, simply to get the best total cost-performance and responsiveness? Would it be best to only run supers at these above-average speeds where they have the best performance/price and provide truly super speed not achievable on the mini-supers? I met one user at Illinois who described his Alliant as operating at 3/4 the speed of the XMP.

Do you think our centers should provide parallel processing service, especially in preparation for the larger mP's such as the ETA 10, YMP, and Cray 3 in order to get some insight so that the large mP's provide the proper environments? We have capacity today to permit this form of use.

## ABOUT PAPER MACHINES AND THE RATIONALE FOR THE MACHINES

I chose the particular machines because the manufacturers gave me the data on them and allowed me to put them in a public document. Also, I considered them to be the leading edge and most relevant machines for our computing environment.

The future SUNs were there because Bill Joy saw the slide and put the numbers on it and agreed to let me use them. I will continue to remind him of them, as they are very aggressive. I wanted to encourage SUN to meet them, because it will provide enormous benefit to the community if they or any other company are successful. (My responsibility is to encourage computing of all types.)

The ETA 10 is included because of NSF's commitment, and because it is a "real machine" to us which we spend time watching and waiting for. NSF allowed it to be bid as a real machine.

The performance numbers are ETA's; I quoted them without comment. Various people have commented on the optimism of their performance numbers. I did not use future machines by Cray or IBM in the document, nor would I ever expect to, because these manufacturers simply don't make future commitments unless they have something running. I hope the ETA situation and numbers are an anomaly in behavior from a bygone era.

I like your slide on the formulation of the probabilities of various future machines. Let me encourage you to tune it up with time varying functions. For example, the probability of shipping at some future time increases up to a point, usually at a reduced performance level. At this point, the probability of ship is reduced as a startup runs out of money, some flaw(s) is discovered (e.g., Trilogy), or in a large company, the people tire and stop the project.

I have always subscribed to Worlton's Rules for Paper Machines. As a former machine builder and manager of groups proposing new machines, a few other rules are needed, of course, if you are to ever let a group propose a machine. The rules of machine proposers who build state-of-the-art machines always stretch the imagination. Ideally, such machines should never be sold to customers. I think we want to tend toward a policy of buying benchmarkable computers. A major part of the rationale for the NSF program was to support the U.S. supercomputer effort, and we are not alone among Federal agencies. I believe NSF was right to obtain the ETA 10 as part of the program despite the risks.

## THE ROLE OF THE NSF SUPERCOMPUTER CENTERS

**Peak Power.** I don't believe we are changing the role of the centers. John Connolly's long range plan has quite a clear statement about their primary role of providing service (including hardware and generic application software) and training. The ASC advisory committee has been concerned about this in light of modest budget growth in FY87. The budget is based on center performance, networking plans and needs, actual use and the congressional appropriation for NSF. In a recent meeting with members of the advisory committee, I agreed that "**the NSF Centers should exist with the highest peak (computing) power -- forever**". You will see this on my board if you visit me. This commitment does not affect our obligation to see that public funds are used as effectively as possible.

**University Centers and Scientific/Engineering Power.** I disagree about your assessment. In fact, university computation centers do manage supercomputers. Note, there are 40 IBM 3090/200's installed now, several XMP's are installed and on order, and I believe the program has stimulated use and ownership already. I think that if we are going to solve the problem of training people in new use of advanced scientific computers, universities must own and operate them. Similarly, we are gathering data on the operational aspects of the centers to serve as a guide for other university centers. I want to include the National Labs in this



comparative data on cost, performance, and service delivered. Could you provide me with such a report on the operation of your center at LASL.

Today, I see the volume of use of the NSF centers by the general scientific community limited by education and imagination (due to education). I don't see the centers addressing this very rapidly as they only have formal courses for a few hundred per year. I would like to see wide-scale teaching of vectorization and parallelization techniques by engineering, science, and computer science departments in the hand crafted way we do now at our centers, at traditional computation centers, and perhaps via smaller computers which are all evolving to have vector capability (if there's any way to buy them). Texts in computing and applications are in dire need of updating. I would like to measure the progress both in terms of the change in degree of involvement by the scientific community and by the final scientific/engineering results. Unless more users get involved and the community is enlarged, I think the impact of the program will be much less than what it should be or that we promised.

**New Machines.** The purpose of the paper was to ask the centers about their roles. For example, they have in fact been beta sites for manufacturer-provided machines including Alliant, IBM and SCS, and I had been encouraging them to continue this role of trying to understand the issue of cost-effectiveness my paper raises, subject to the constraint of providing service and having the latest and fastest machines. It is desirable to understand how smaller machines could act as problem feeders for a larger base despite the fact that the centers program discriminates against this level. I suspect users who don't have unlimited free resources will actually choose smaller machines. Do you think the centers should be involved in smaller machines in any way?

It is now pretty clear to me that the centers should not be involved in highly experimental machines (e.g., the Connection Machine, or GF11), even though I expect such machines to have the greatest power for a number of problems in the future. These machines are not general purpose now, and the centers already have too much to do already. Do you have any ideas how these machines will be funded or evaluated for scientific use?

**Standards.** The NSF centers could become involved in the standardization process. History has shown that standards are the key to sharing software and databases, and building a large user base. I see having a compatible uniform system across the hierarchy of machines as the key to extending access to users at workstations and remote sites that require scientific/engineering computing. I don't believe NSF should be funding the development and maintenance of an operating system, especially one that is at variance with forthcoming Federal Information Processing Standards and is outside the standards process, operating under its own laws. NSF simply doesn't have the funds to participate in building vanity operating systems (e.g. CTSS, and its successor) and language dialects (e.g. LANL Fortran) that characterize the National Labs' environments. Our centers must build on the past work of others, especially the manufacturers, not reinvent systems that have to be maintained. I am delighted that Cray has introduced Unicos, and a number of sites are running it (Berkeley, Bell Labs, Cray, and several in Europe).

I don't think the centers should be engaged in any systems development, but rather should promote the development of a high quality, standards-compatible environment by knowing what they need--communicating with their suppliers and participating/leading the standards process. This includes all the aspects of a standard environment: networking, the programming environments (especially one that dramatically improves Fortran), human interfaces, graphics, parallel processing, etc. I believe the centers have the resources and obligations to provide a common environment that is compatible with the relevant Federal Information Processing Standards such as POSIX, TCP/IP, etc.

At a recent meeting, all of the center directors agreed to provide the POSIX (UNIX) environment in the future. Allen Newell, who chaired the meeting saw this as a significant event. I am anxious to see a schedule and milestones.

## CAN THE CENTERS PROGRAM ELIMINATE SMALLER COMPUTERS?

Stuart Rice, a chemist and supercomputer user at the University of Chicago, and recent National Science Board member, raised the question of the funding mechanism for the mini-supers because of the high share of NSF funds for scientific and engineering computing that are directed into the centers. In a talk at the IBM ACIS Forum on the Physical Sciences, he described the hierarchy and stated four basic concerns (which I agree with) as follows:

"I have in mind a networked system ... graphics workstations and local supporting intermediate computer and ultimately connects to a supercomputer with provisions of special devices ...

1. Distribution of computer resources is distorted by the use of "funny money" ... cash and credit ... Workstations come from grants, supers are "free" ... intermediate machines are indispensable and current funding patterns will have to change if ...
2. Dramatic advances in hardware haven't been matched by advances in algorithms and operating systems. ... parallelism is "chicken and egg"
3. The scientific community has become rather inflexible with respect to use of operating systems ... don't use particular machine features...
4. ... the scientific community has not been as imaginative as it might in thinking about the uses of computation in research."

I have been developing a census of scientific computers. It shows that most scientific computing is done on "departmental" machines, and many worry that the centers program could act to drive out this style of computing, independent of its merit. Note the distribution of computing power, given in Cray 1 equivalents:

Supers:	60	from 30 processors (counts a 4 Pc XMP as 4)
	80	from 40 - IBM 30390's
IBM Mainframes:	14	from 80 - 3080's
Mini-supers:	15	from 30 Alliant, Convex, and SCS
Super-minis:	75	2650- VAX 7XX ... 250 VAX 86XX, and 250 IBM 4381)
Workstations:	16	from 1600 microVAX, plus
	??	Apollo's, IBM RT's, and SUNS
large FPS	20	from 50, ≥164's

The paper, and above count, and power equivalence doesn't mean to imply that computing power can be used interchangeably on problems any more than one would imply common utility in evaluating electrical power capacity across large power stations, home stand-by power generators, and batteries, etc. In a similar fashion, given that your paper on computer export control stated that it takes about 8,000 hours on a supercomputer to design a nuclear weapon, we might conclude: the Russians don't have any weapons, or that it takes them much longer to design and theirs are obsolete, or they use better algorithms, etc.

As you know, many start-up and traditional companies are pursuing new, scientific machines that should be very cost-effective. I'm anxious to make the comparison with the next Crays. I've already had one inquiry from a reporter as to whether the centers program might kill companies building small computers. Since NSF discriminates against minis (super-minis, and mini-supers) the many users of smaller and experimental machines might make similar arguments to restore funding for departmental/group level machines, just as the original arguments were made by the folks who needed supercomputers.

## INTERACTION WITH THE PROGRAM

To the best of our knowledge, no one in the Directorate is involved in the Iranian-Contra crisis. You have a copy of the document I wrote because I asked for comments on it. No use is being

made of it other than to solicit comments regarding the issues it raised, including this one on small machine discrimination.

NSF wouldn't consider operating without consensus. All of us in the Directorate spend a great deal of time and effort interacting with users, Centers Directors, and the Advisory Committee on these and other issues. The consensus which launched the ASC program never was restricted to, as you imply, specific facilities. We intend to keep the program at the frontier. Moreover, we have the responsibility to evaluate it and to manage it according to its progress and the emerging scientific computing environment (such as products from IBM, the whole class of mini-supers, specialized mini-supers, and high performance experimental machines that achieve super performance) not envisioned in the original proposal. That is critical to maintaining a broad, foundation-wide consensus and not permitting it to degenerate into what can be seen as consensus of a very limited "constituency" that does not advance science or engineering widely, nor communicate with the industrial community.

Thanks for your comments on the paper, I look forward to working with you and your colleagues at the National Labs for continued understanding about the performance and operation of current and future computers, and ultimately to the creation of a better scientific computing environment.

Sincerely,

Gordon Bell  
Assistant Director

cc E. Bloch  
J. Connolly  
R. Ewald  
Derek Robb

Note COPY

LETTER TO CENTER DIRECTORS

August 29, 1986

Dear

I look forward to meeting with you at the SC Director's meeting on October 7, 1986, at Cornell University. I am particularly interested in discussing several issues which will be crucial to the health of the Foundation's support for the supercomputer research program over the next several years.

Thus far, NSF expects to spend \$311 million in supercomputers over the next 5 years. An additional \$200 million comes from other sources to achieve the originally envisioned program. Expenditures at this level raise strong expectations about the scientific efficacy of the Centers and the nature of their operations. As we gain operational experience with the Centers, it will be increasingly important to document various aspects of the program and to devise policy to maximize its usefulness.

Based on my observations and information available to OASC, I believe that we need to take major steps in improving information about the following which would lead to higher quality, efficiency, and effectiveness:

1. Service provision: In addition to time allocations we need data which is comparable across centers, and useful for analyses of the cost/efficiency of alternative computational resources (within and beyond the centers);
2. Service use: Apart from allocation, we need to identify the scope, nature, density, (within disciplines) and adequacy of service from the researchers' point of view. It is especially important to characterize the nature of the demand and how it varies over time, and;

3. Scientific and engineering output: The research community should be concerned with improvements in scientific knowledge, not computer time. If the centers are effective, we should observe a significant change in research output by the user base. Although this final output is hard to measure, we should start to see results that are unique to the Centers' program. What are they? How might we begin to get at the merit of the research and the likelihood of breakthroughs through computation?

In addition to information, I believe that we need to become far more concerned with the overall operation of the program. There are major opportunities through policies about the following:

- (1) Standardized software libraries, and common operating systems and user environments;
- (2) Non-duplicative disciplinary emphases, specialization (specialized data resource requirements) and improved access;
- (3) The desired end user environment and workstations, software for access, graphics, advanced utilities, etc;
- (4) Shared use of special facilities, unique machines and software, and;
- (5) Optimal time allocation schemes based on actual cost efficiency of alternative computing resources (at or outside of the Centers)

My overall impression is that the Centers have done a marvelous job of getting started. In the face of the current budget environment in Washington, and rapidly changing scientific computing technology, it will be necessary to devote equal energy to operational and evaluation issues. I look forward to discussing these things with you, and invite your suggestions and reactions.

Sincerely,



C. Gordon Bell

cc: John Connolly, OASC  
Frank Stillinger



18 November 1986

Professor Michael Dertouzos  
Laboratory for Computer Science  
545 Technology Square  
Cambridge, Massachusetts 02139

Dear Mike,

I enjoyed the meeting with you, Gerry Sussman, and Piet Hut on Saturday regarding the proposal to build a Computational Observatory for Galaxy-Galaxy Scattering experiments. I believe that Computer Science and Mathematics should consider this work on Computational Science aimed at the "third paradigm of science: simulation" to be perhaps our highest priority at this point in our collective development. The building of these experimental laboratories that do virtually all of the experimentation through computation also needs to be an equally high priority within each of the scientific disciplines we would hope to "revolutionize". Chemistry, especially molecular synthesis supporting bio-technology, is probably at the top of the list in terms of need because of the complexity and engineering nature of the problem.

NCAR, I believe, has operated in this mode for a decade by providing computer resources (cycles, programs, and databases) that a common community of atmospheric and earth scientists at all universities interact with to carry our simulation with experimental verification. This model of distributed, science surrounding a common set of computational resources is clearly the next step in the evolution of virtually all science, to both supplement and in some cases supplant experimental science. Let me urge you to use some mechanism such as workshops and seminars to explore this. Ken Wilson, Nobel Laureate, at Cornell and a researcher in Electronic Structures is probably the most lucid and outspoken on this viewpoint. He is the Director of the Cornell Theory Center, one of NSF's centers. They have the 3090. Let me urge you to get him to talk to the entire scientific and engineering community at MIT.

This should be the right time to undertake this work because the machines are available, can be made available, or should be the subject of construction and research by the computing community. NSF's five National Centers are now operating 3 Cray 48's, a CDC 205, and an IBM 3090/400 with vectors. NCAR, which NSF operates, has just installed a Cray 48. I believe you should use computer time at our centers to start experiments in several areas in order to verify how amenable the various areas of science are to

simulation, and to estimate computational needs for various problems.

The Digital Orrery by Sussman, which computes at 1/2-1 Cray 1/S for orbital calculations is an excellent testimony to the notion that specialized machines can be used for scientific computational experiments. IBM is building the 11 Gigaflop machine for QCD calculations to compute the mass of the proton. Columbia/Brookhaven are also building a similar machine. CMU's Chess machine is another example, and we have numerous simulation machines for engineering. Whether you need a special machine or not would be a product of this next exploratory phase.

Independent of whether you need or build specialized computers for these various laboratories, I believe it is critical for MIT to acquire its own, general purpose scientific computer, i.e. a supercomputer! I have been surprised to find that MIT faculty are making only minimal use of NSF's supercomputer centers. Thus, I see MIT as lagging in this form of science.

Let me urge you to mobilize several of the scientific disciplines to establish a laboratory (with its own computer) capable of combining the discipline, mathematics, algorithms, and appropriate parts of computer science, including artificial intelligence. I believe MIT is uniquely qualified to take on this revolution in science because of its size and strength in both engineering and science. These "Computational Laboratories" simply require significant engineering beyond anything computer scientists builds today. Thus a by-product is likely to be a much greater gain in software engineering than anything that comes from researching software engineering per se.

MIT has played a very important role in computing with Whirlwind in the 50's, timesharing in the 60's, and its contribution to artificial intelligence because it actually built large scale systems. Berkley and Carnegie-Mellon are the only other two places which could organize such an effort, I believe. Within a year or two, it could be in a position to lead this revolution by establishing a laboratory center for scientific research through simulation. Let me urge MIT to proceed.

Again, I enjoyed the stimulation interaction with you, Gerry and Piet on Saturday. As you can see above, I share your vision. Now just follow it.

Sincerely,

Gordon Bell  
Assistant Director

cc Erich Bloch, Director  
John Connolly, Division Director, Advanced Scientific  
Computing  
Paul Gray  
Rich Nicholson, Assistant Director, Mathematics and Physical  
Sciences

18 November 1986

Dr. Paul Gray, President  
Massachusetts Institute of Technology  
77 Massachusetts Avenue  
Cambridge, Massachusetts 02139

Dear Dr. Gray:

I just wrote to Mike Dertouzos urging him to take on the revolution in science that we believe will come with the availability of high speed scientific computers. A copy of the letter is attached.

MIT has not been involved in this revolution to any degree so far, I believe. Furthermore, MIT is not a very large user of NSF's Supercomputer centers, nor does it have significant, powerful local scientific computing capabilities.

Given the fine relationship you have both with IBM and NEC, let me urge you to urge them to provide you with one or more of their scientific computers as components for the laboratory or center that is described in the letter to Mike. This would allow the laboratory to get started now.

Sincerely,

Gordon Bell  
Assistant Director

cc Michael Dertouzos

Dear Joe:

In discussing your letter with Gordon Bell, Assistant Director of Computers and Information Science and Engineering (CISE), he made the following comments. They pretty much confirm your observations, and in addition, he raised an issue about the future of Cornell vis a vis upgrading the 3090/400 to a 3090/600, and Cornell's commitment to production, parallel processing.

The Cornell Center was initially established as a composite entity with at least one major goal of supplying supercomputer power to the entire scientific and engineering community through the Production Supercomputer Facility. This implies distributing at least half the power outside of Cornell. In addition, several other independent and orthogonal activities were part of the Center: Don Greenberg's graphics research, Ken's work on building experimental parallel computers, and Gibbs.

The Advanced Scientific Computing Program is now a fantastic success, with 6000 supercomputer users and 2000 research projects at 200 universities. Furthermore, many universities are buying their own supercomputers. It has become increasingly clear that:

- the *demand for quality* computational power from the users,
- the *necessity* to keep the latest equipment at the centers,
- the *need* to link centers and users via a high performance network,
- a *desire* to establish grand challenges (Ken's words) formed around teams of computational scientists formed in every discipline, and
- the *possibility* of doing exciting research into parallel computers, visualizing results, new algorithms, etc.--

all of which can absorb large resources, the overall program must be segmented into different sub-programs which are managed and funded independently. The program's goal is now simply to supply quality, computational power (and training) with the highest performance machines at the five National Centers. This is what the congress, NSF, and the users want and expect from the program.

CISE has also established a separate program (see attached letter) addressing a broad range of important topics in advanced scientific computing, including: new technologies such as parallel computing, visualization, and the formation of computational science teams aimed at specific disciplines. This research program is operated on an independent basis. In addition, we have a foundation-wide, discipline specific program on Computational Science and Engineering (see attached) that's distributed through every discipline.

In looking at the trajectory of the Cornell center, Gordon saw no synergy among the activities, and the scope of a substantial part of this work is well outside of the charter of the ASC program.

Don's work on Graphics continues to be exciting, but has nothing to do with either supercomputing or parallelism. In fact, at Gordon's urging, the FPS machine in Don's Center was apparently moved back to the computing center so that it could be used in a parallel processing fashion by all users. This was designed to provide much better service, such as parallel processing, for all users, including Don. Don's graphics work can be funded, as it has in the past, in various places throughout the foundation and other agencies. If it is focused on the visualization initiative, then he can apply to the program described in the attached letter.

The Gibbs project proposal painted an exciting scenario, but was abandoned because you lacked the manpower. In restarting it, the pressures on resources at your center, clearly put it below your cut line. Since Gibbs continues to be an exciting idea, it should have no trouble in being funded here or elsewhere.

Before the 3090/400 was installed, Cornell came with a proposal to upgrade the FPS machines and install more at a significant cost to NSF. Gordon apparently discouraged this because of the efficacy, lack of generality, and reproducibility of the system. Today, no one wants to use these machines and Cornell is considering getting rid of them because of poor reliability, difficulty of use, and high operating cost. Fortunately, Cornell is beginning to see the 3090 as an important, parallel machine, and attend to the need to supply computational power to the community.

Cornell is spending money doing conventional systems programming on compilers and operating systems for the Floating Point Systems T-series. The T-series was supported by the Federal Government, through DARPA, at a cost of \$10 Million at Ken's urging, even though Cornell refused the "poriky" machines. Cornell (NSF) is still part of this ill-fated effort. Gordon's feels that the T-series is still inadequate for virtually every task, even though all but the switch have been redesigned. He claims the switch will have to be redesigned if users operate the computer for anything other than a limited set of applications, such as Monte Carlo calculations. Again, NSF is not happy about funding conventional software development that should be supplied by a manufacturer. This work will not be considered to be part of the centers program in the future. It would seem that DARPA would be an ideal place to get funding for this work, since they had to fund the T-series. Also, CISE has two programs which might fund this basic work.

Finally, networking is a separate division of CISE charged with providing a network for supercomputer users to access centers. This was done for a number of reasons, including an organizational one within CISE, aimed at making the network operate! The lack of good networking has been a serious impediment to using the centers, especially Cornell. For example, many of your users come in via 800 dial-up lines, not the network. Given the critical

nature of networking, Steve Wolff, and Gordon are both spending a significant amount of time working in this area. Gordon heads the Federal Co-ordinating Committee on Science and Technology, for Network Access of Supercomputers. Ken Wilson and Allison Brown have been part of a study group, which recently met in San Diego, working on this critical issue. NSF views networking for both the supercomputer access and for the interchange of scientific knowledge as a critical thrust.

Recent meetings with your center seemed to indicate that it is heading toward providing supercomputer service based on the IBM 3090/400. Furthermore, there appears to be an excellent relationship between all parts of IBM and Cornell. IBM is excited about the possibility of having Cornell being its leading edge site for parallelism research. We see this as a great possible opportunity for Cornell, provided it can first muster and then focus its resources.

IBM recently visited NSF concerning an upgrade of the 3090/400 to a 600, and whether we felt the Cornell center contract would be renewed. We also discussed funding for Cornell being the test site for a second more loosely connected 600.

Although the renewal depends on the future site reviews, NSF has been encouraging Cornell to get in a position to pass a review as a national facility. It is our desire to see Cornell become a healthy center fulfilling its role. In addition, Cornell is in a unique position to provide the greatest supercomputing power on a continuous basis using parallelism, if the 600 upgrade takes place. This later role takes a commitment NSF doesn't see at this time. This would mean focusing resources to understand, explore, and support parallel processing on the 600. Such an effort would include systems programmers such as those we are apparently supporting on the T-series, computer scientists, and computational scientists with discipline specific problems. Cornell is in a unique position to do this work, but other universities who have 3090/400's could also carry out the work. NSF and IBM both believe the work is critical and want to proceed with it.

Are you interested in working on this difficult, main line focus on parallelism based on the 3090 at this time, or do you prefer to explore the alternatives which are unlikely to yield significant, general purpose supercomputer power to your broad user basis?

Although Erich, Mary Clutter and I are the only part of NSF charged with an entire overview of science, including computational science in every discipline, the CISE directorate has a very broad overview of computing., including some programs in Computational Science and Engineering. Gordon Bell and John Connolly are anxious to discuss any of the above issues you feel are necessary. In particular, the 600 upgrade sounds like an especially critical issue that should be settled immediately.

Sincerely,

John Moore

CC:

Ken Wilson  
Gordon Bell  
John Connolly

*Also like  
them for  
centers*

Subject: Thoughts and Concerns on Centers Operations  
To: Supercomputer Center Directors  
From: Gordon Bell  
Date: 25 August 1986

The Centers have done an extraordinary job of proposing and establishing large, computation facilities, building networks, acquiring software, and building a generally happy user base. By all accounts, a number of scientists and engineers appear to be utilizing the centers effectively. OASC reported to the National Science Board in August, and received approvals to operate the San Diego, Illinois, and von Neumann centers for periods of 3, 3, and 1 year respectively. At the Board meeting, updates were given of the various centers in order to provide some understanding of the general status of the operation.

Looking at the operational data, studying the site visit reports and user questionnaires, visiting two of the centers, and talking with various users raised many concerns in my own mind. I think it is time to review just where we are in the program.

The report to the board was based on sketchy and inconsistent data from you which I hope is just poor reporting. Hopefully, it is not an indication of the operation. Our goals have to be providing very high quality service. I would like to see these areas addressed at the Directors meeting and then operationally:

1. Services supplied, including cost and efficiency data- We simply have a poor understanding of the services being supplied, based on many of the reports. Rather than trying to provide a common data base at NSF, it would be preferable to get quarterly reports, in some common format, on the services the centers are providing. These reports should provide cost and efficiency data for services and compare this with other alternative forms of computation, including: resource centers, single discipline centers (e.g. NCCAR), public utilities, local mini-supers and workstations, and specialized computers.
2. Services used and the community of users- NSF is especially concerned with fairness. The centers were established to provide service to the entire community, not just the local university, the consortia, or a single intellectual community. The concern is across: geography, discipline, and user size (especially the beginner). While the question of fairness is not significant now while we have idle time, it will become a major issue as the demand increases. (If the demand doesn't increase, the centers will have been poorly conceived.) On the other hand, we observed that there is idle time in the centers- this is a poor utilization of our national resources and this problem must be solved immediately. The atmospheric scientists are making good use, because they have been long time supercomputer users and programs always ready to soak up any amount of time.

The data on the use shows institutional locality, and a very



small user base. A few hundred significant, mostly local users out of 15,000 NSF grants are utilizing the centers. Maybe this is the way it should be. My own feeling is that unless a user is getting an hour of time per day, on the average, he'd be better off on a smaller machine. This would mean that a maximum of 4 x 160 "average" users would be at the centers.

3. Scientific and engineering output-The community should be concerned with final scientific knowledge, not computer time. If the centers are effective, we should observe a significant change in research output by the meager user base. Although this final output is hard to measure, we should start to see results that are unique to the centers program. What are they? (Our Nobel Prizes cost us in the neighborhood of \$12 billion per year, so any incremental prizes are a bargain for the \$0.5B we'll spend over the next decade at the centers.) Just what are we doing to look at the merit of the research and the likelihood of breakthroughs through computation?

4. Standards, common library software, and ease of use- Since three of the centers operate Cray XMP's, we have the opportunity to use a common operating system, common libraries, etc. so that users can operate on any center, including computers outside the centers. Given the large number of systems that use Unix as an operating system, I believe we should adopt this as the standard operating system as soon as possible. This will allow user compatibility with workstations for graphics and local computation, and the mini-supercomputers such as Alliant, Convex and SCS, all of which provide computation on an equally cost-effective basis as the Cray. Similarly, the Cray 2, 3, and ETA are all providing UNIX. What are the libraries we want across the centers?

5. Programs and data for communities- It would seem that certain communities need common data and/or programs along the lines as NCAR is operated. This would argue for both having centers specialized when large datasets are involved, but certainly common operating systems in order to support use across a broad community without reprogramming. What is or will happen here?

6. Workstation and graphic support- Given a single user may want to utilize several facilities, it is essential to have a common set of local utilities for graphics and computer access. Certainly, it seems necessary and desirable to support all the graphics of all of the centers, provided a user will be accessing multiple centers.

7. Utilizing special facilities- Today, Cornell has two highly specialized facilities from Floating Point Systems capable of supplying a very large amount of computation. These require programs to exploit their unique capabilities.

While there are certainly more questions and issues involved in providing computational service at the level our users should demand, the above should let us start to focus on the question

of providing high quality output and easier access.

I look forward to interacting with you at the Cornell meeting.

October 23, 1986

^F1^

^F2^

^F3^

Dear ^F4^:

Enclosed is a letter to the centers directors introducing issues that I believe are critical to the measurement, operation, and future of the centers. Some of these issues were discussed at the Centers Directors meeting October 6 and 7.

You have a paper (also distributed at the meeting) outlining a set of issues that I derived by looking at today's scientific computers. It shows that machines exist in a hierarchy, permitting users to compute in a variety of styles depending on various factors such as networks, graphics, standards, local programs/databases, and economics. Since the data show that all styles of computing will exist, it should be our goal to provide the most productive environment for scientists and engineers today, and in the future. This strongly favors having a "standardized" environment so a user can chose any hardware at a given class and level (e.g. supercomputer, mainframe, super-mini used as a front-end or back-end to a supercomputer, mini-super, graphics workstation, or personal computer), and be able to migrate work among the levels as needs and machines evolve.

② Michael Levine of Carnegie-Mellon University is convening a conference at CMU to examine the issues in supporting such a common environment which includes: network, front-end, graphics, programming environment, and user interface, etc.

① In addition to the issues described in the letter and attached report, I am quite concerned about the limited availability, use, and growth in use of the centers. Ken Wilson has convinced me that computation, characterized by today's supercomputers is the next revolution in science. I believe the same can be said for many parts of engineering. With the exception of San Diego, the use seems to me to be quite local and limited. I believe the user base is not expanding as rapidly as is necessary to justify the 1/4 billion dollars that will be spent for this program. We must understand just what is happening regarding training and use. I think the Advisory Panel must address this issue as their highest priority!

*Also see W. Wilson's survey of  
the current state of scientific computing*

I am committed to help you make the centers the leading places for scientific and engineering research through a new paradigm of computation. The program has fine support from Congress, other NSF Directorates and manufacturers. For example, Alliant, IBM, and Scientific Computer Systems have provided their first machines at no charge to centers for evaluation. I am encouraging other hardware and software suppliers to follow their precedents.

You are meeting on November 20-21. I hope you can start to look at what's been accomplished during the first three years of the program, and then address critical issues and opportunities.

Sincerely,

Gordon Bell  
Assistant Director

Budget

CPT

M E M O R A N D U M

April 8, 1987

FROM: AD, CISE  
TO: DIRECTOR, NSF

TOPIC: Advanced Scientific Computing FY88 Budget Issues

This is to detail budgeting considerations for ASC centers, and to illustrate the need to work out an NSF budgeting strategy recognizing ASC (and network operations) as facilities serving all of science and engineering. My concerns are as follows:

1. ASC requires an additional \$3.5m to \$9.5m added to the FY88 budget. The FY1988 Congressional budget request for the centers component of ASC is \$45m. This is \$3.5m less than required to meet current cooperative agreements and, in addition, to renew the agreement with Cornell. At the \$45m level, the program may be forced to reduce one or more centers to levels too low to justify continued support. A maintenance level for the program (allowing needed staffing additions) would require \$5m. This will not guarantee the leadership role of the NSF centers. To achieve leadership requires \$9.5m. The increments, the staffing additions and equipment upgrades purchased, and priorities across the range are shown in the attached chart. As reference, the annual NCAR budget of \$18m is also noted.
2. Additional funds for ASC must be drawn from the NSF budget rather than from the CISE component. Reductions from your initial assignment, required to meet the OMB mark, were made entirely within CISE. In making them, I preserved the priorities developed in the first CISE LRP, proportionately reducing all CISE activities. This reduced the aggregate percentage growth of the relatively small CISE research programs to substantially less than comparable activities in Engineering. Further tradeoff of the ASC activity with the rest of CISE is inappropriate, given the broad goals of ASC and the small base of CISE.
3. Several things contributed to this situation: the success and popularity of the ASC activity, the opportunistic nature of computer facilities programs, the requirement for unplanned but essential upgrading to stay at the state of the art, and the structural change in the NSF budget. As the attached "Leverage" chart shows, NSF cannot depend on uncontracted industrial sources of leverage, but accept them opportunistically. Given these reasons for strong agency support of the ASC program, a correction is needed to the budgeting philosophy and process. ASC (and Networking operations) must be detached, made separate "NSF" lines not part of the CISE budget, while remaining part of the organization.

ATTACHMENTS

ACS BUDGET ANALYSIS  
(\$millions)

	FY88	FY89	FY90
Current Centers	\$45.0	----	----
Cooperative Agreement (+Cornell, +JVNC/ETA 10)	48.5	51.9	54.6
+minimum needed staff additions PRIORITY = 1	1.5	1.7	1.7
+3090/600 at Cornell (4 yr payoff) PRIORITY = 2	2.5	2.5	2.5
+fix memory problem and upgrade SDSC PRIORITY = 3	1.6 ---	1.0 3.8	1.0 3.8
+NCSA upgrade correction PRIORITY = 4	---	1.5	1.5
+Pitt upgrade correction PRIORITY = 5	---	.8	.8
TOTAL	\$53.7	\$63.2	\$63.9

FOR COMPARISON:

1. Rotar model of need..... (5 yr amort.)	\$54.5	\$63.2	\$68.9
(3 yr amort.)	\$59.0	\$73.0	\$78.9
2. NCAR annual	\$18.0	\$18.0	\$18.0

### HISTORY OF FUNDING SOURCES

	SDSC			ILLINOIS			JVNC			CORNELL			PITT	
	85	86	87	85	86	87	85	86	87	85	86	87	86	87
NSF	4.9	9.5	12.1	5.1	7.2	8.7	8.8	5.6	8.5	5.0	5.4	5.4	2.0	4.0
STATE	1.0	1.0	1.0	1.0	2.0	3.0	0.6	1.5	0.6	0.8	0.8	0.8	1.9	1.0
UNIV.	1.2	1.2	1.2	1.3	1.5	1.0	---	---	---	1.8	0.7	0.7	3.6	2.7
TOTAL	7.1	11.7	14.3	7.4	10.7	12.7	9.4	7.1	9.1	7.6	6.9	6.9	7.5	8.6
% NSF	69%	81%	85%	69%	67%	69%	94%	79%	93%	66%	78%	78%	27%	47%
VENDOR	0.5	0.7	0.7	1.5	1.5	1.5	1.8	1.8	1.8	5.7	16.5	8.6	0.5	0.7
TOTAL	7.6	12.4	15.0	8.9	12.2	14.2	11.2	8.9	10.9	13.3	23.4	15.5	8.0	9.3
% NSF	64%	77%	81%	57%	59%	61%	79%	63%	78%	38%	23%	35%	25%	43%
INDUST.	0.4	1.9	1.6	0	1.6	1.0	0	2.5	0.2	0.3	0	0.2	0	0.5
TOTAL	8.0	14.3	16.6	8.9	13.8	15.2	11.2	11.4	11.1	13.6	23.4	15.7	8.0	9.8
% NSF	61%	66%	73%	57%	52%	57%	79%	49%	77%	37%	23%	34%	25%	41%
CONSORTIA	4.0	4.0	4.0	---	---	---	2.5	2.5	2.5	---	---	---	---	---
TOTAL	12.0	18.3	20.6	8.9	13.8	15.2	13.7	13.9	13.6	13.6	23.4	15.7	8.0	9.8
% NSF	41%	52%	59%	57%	52%	57%	64%	40%	63%	37%	23%	34%	25%	41%

April 14, 1987

Professor Niklaus Wirth  
ETH-Zentrum  
Institute for Informatics  
Zurich, Switzerland 8092

Dear Niklaus:

Enclosed is a copy of the memo which I wrote raising various questions about the operation of our supercomputer centers. I have another memo in progress which compares what I think are the pro's and con's of a "centers" approach with using a plethora of group and departmental, distributed superminicomputers. Basically, it shows that our centers cost NSF about \$11M per year to run, and someone else (state, university, industry) kicks in another \$5M. In our case, I believe an \$11M expenditure would buy anywhere from 50-100 mini-supercomputers, depending on how they are run, and provide 3- 6 times the aggregate power of a center. Furthermore, our centers serve only 5-10% of the scientific user community, and an even smaller part of the engineering users. Finally, by having centers, we are discouraging users to buy their own computers and use our network of supercomputers. (Our networking costs about \$2,000 per user.)

While I believe NSF needs to sponsor centers where the greatest computing power is available for the top few jobs, I don't believe it is an especially good way to supply scientific and engineering computing to the large number of users and potential users. I would somehow like to encourage users to have their own mini-supers either in individual labs or departments with a minimum staff. Small machines should be compatible with the centers where larger jobs could be run. Right now they aren't because our centers can't run Unix.

Since I don't know the complete environment (i.e., availability of supers at service bureaus, labs, etc., and the university structure) in your country, it is probably presumptuous of me to offer a solution. My first reaction is: get as many mini-supers (perhaps all with virtual memory) as you need with UNIX running on them, get compatible UNIX workstations for graphics, etc., build a good LAN environment, and perhaps allocate some funds for a few user to buy supercomputer time externally when their problems demand it. Learn about this kind of computing. Track user demand by encouraging users to buy machines when they need them.



My other reaction is: get a Cray-XMP and run UNIX with a LAN to UNIX workstation. It has lots of software and the community is producing more for it all the time. It may solve a few more large problems, provided the "system" will allocate the time to a few, very large users.

Sincerely,

Gordon Bell  
Assistant Director

Enclosure



NATIONAL SCIENCE FOUNDATION  
WASHINGTON, D.C. 20550

OFFICE OF THE  
ASSISTANT DIRECTOR  
FOR COMPUTER AND INFORMATION  
SCIENCE AND ENGINEERING

April 2, 1987

Professor Edward A. Feigenbaum  
Computer Science Department  
Stanford University  
Stanford, California 94305

Dear Ed:

Your question, roughly: "If AI \$ is so great as measured by numerous companies and organizations making money and doing new computer applications, why aren't the AI companies making money?", stimulated the following thoughts (and call to action).

In general, the AI suppliers are simply designing and building new languages to build expert systems. I don't see them as AI companies at all, and rule-based programming is merely another programming technique that should be acculterated into traditional computer science.

There are several reasons why "language" suppliers, in general, have never existed to a large extent and made much money.

Language construction is a highly creative engineering intensive business. "Garage shops" appear to be the best place to develop them. Look at the language business today. A tiny company called Green Hills develops most C's, Fortran's, and Pascal's. Ryan-McFarland and Microfocus supply Cobols. Several companies are going after the potentially lucrative ADA market that accompanies the waste surrounding military procurement. Gold Hills, Lucid and Franz are building LISPS.

Large companies who have traditionally been good at languages oftentime build their own (such as DEC and IBM).

The "systems" companies are inherently the "best" place to sell, distribute, train, and support products--particularly as "expert systems" building languages begin to look like commodities and traditional programming. This will mean "bundling" and lower prices for the "tools", i.e., languages to build expert systems.

The early establishment of the three companies to supply these in the mainframe, mini, and workstation market (Teknowledge, Intellicorp and Inference Corp) was an anomaly and needed at a time when users required training of the programming methodology. I suspect Teknowledge could be doing the best because it supplies training, runs on a number of computers, is therefore somewhat of an integrater to a customer, and does special systems. This enables them to understand end user needs.

In the PC world there are many companies building and selling these languages because the task won't stand the support (at the low prices). Ultimately, I'd expect a distributon company like Ashton-Tate, Electronic Arts, or Microsoft to be the best distribution channel--although an established company like Gold Hills could get the market.

Also note you have too many, high overhead, high product cost suppliers, including DEC with OPS5, fragmenting and confusing the market by not having a standard language. This means "expert systems" can't be easily taught in traditional CS curricula and other places in the same fashion as Basic, Cobol, Fortran, Pascal, etc. There's no reason this can't be taught to high school students. Users also delay buying.

The place to make money, using the "traditional" computing market as an analogical model is in "applying" the language to solve real problems. The lack of a standard is probably impeding this, even though you've found lots of real use of the existing competitive tools. The situation is akin to programming business and scientific applications prior to Cobol and Fortran.

Two places apply tools: end users (so far these applications are fairly simple) and companies who have expertise and provide tools in vertical or end use, e.g., Syntelligence. For example, today's CAD suppliers use rule-based approaches for some of their work. They make money because they know CAD of logic or VLSI, not because they know expert systems. Expert systems builder languages are only tools to make real end user tools.

The point of this letter is not only to give a simple answer to the question, but to raise a broader question: How can there be more use of rule based systems technology for building expert systems?

I believe the answer to the question is standardization! By standardizing, traditional computer science would embrace the rule-based approach, teach it, and we'd see a much greater

proliferation of these tools. This would let the AI types work on hard problems rather than being plain old systems programmers. Recall that I encouraged Steve Squires at DARPA to get the LISP community to agree on Common LISP as a means of consolidating the fragmented AI community. Using Prolog or OPS5 (lacks a procedural interpreter) and LISP (lacks an inference interpreter) isn't the answer because this forces users to build mechanisms as the AI community did for years before they built systems. They all require more primitives for data and knowledge bases.

My solution to this problem would be to extend the two dominant systems programming languages, C and LISP, to have really high quality inference mechanism. Ultimately ADA would be extended perhaps and this process could be conducted via the DARPA community, and Steve Squires might be the best to sponsor this since he succeeded in doing the same job on common LISP.

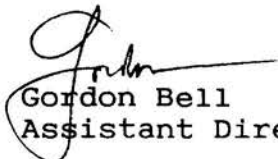
An alternative approach would be to build another parallel language that could be coupled to any traditional language, but I doubt if this is right.

Of course, we could wait till one of the approaches becomes dominant, but this can't happen easily the way I see it.

What do you think of this view?

Please feel free to distribute this.

Sincerely,

  
Gordon Bell  
Assistant Director

cc Sam Fuller, DEC  
Lee Hecht, Tecknowledge  
Tony Slocum, Lucid  
Steve Squires, DARPA  
Y. T. Chien  
Steve Squires

x March 1987

nnnn  
iiiiii  
ssss  
tttttt

Dear Dr. Large, Supercomputer user:

I was responsible for the the first mini- and time-sharing computers and later VAX. Now, I would like to introduce you to Titan, the first Graphics Supercomputer from Ardent Computer. I believe Titan will be the most important scientific and engineering computer since the VAX and the Cray 1. The reasons are simply:

1. Titan provides 10 Megaflops (measured on the 100 x 100, 64-bit Linpack benchmark) of computational power. This is roughly the speed of the Cray 1 or the IBM 3090 processor with vector facility. It has a memory of 1 to 16 million, 64-bit words (8-128 Mbytes), and a full 32-bit virtual memory for large programs.
2. Titan costs less than \$100,000 so it can be easily purchased for a project. This makes it roughly 5-10 times more cost-effective than a supercomputer.
3. By utilizing Titan as part of the "project" team, you get 15-40 hours per week of dedicated, "equivalent" supercomputer time (750-2000 hours per year). By comparison, only a few projects are allocated over 1,000 hours per year at NSF's six national supercomputer centers. By returning to "distributed" computing, whereby a project operates its own computer when necessary, the "average" 50 hour project at a center gets 15-40 times more computing.
4. Titan provides a "free" Cray processor's worth of power for visualization. By being able to render pictures via Dore' our Dynamic Object-oriented Rendering Environment at a real time rate of 400,000 shaded polygons per second, we believe new insight and applications will result. Titan provides interactive, not batch graphics. Turn-around is instantaneous, including video production.
5. Titan is fully compatible with your existing computing environments. It provides VMS and Cray Fortran, C, Unix, and Ethernet. Ardent supports and encourages standards. We are licensing Dore' to be utilized on all computers on the same basis as Unix.
6. A substantial number of applications level programs exist and are being ported now to run on Titan including: Matlab, NAG, Nastran, Gaussian 86, y, z. ....

In summary, I feel we've built a tool that will cause again revolutionize science and engineering. Let me urge you to order one today. If you don't like it we'll refund your money (you'll be the first). A price list, using the NSF discount schedule is included.

Alternatively let me know if you want more information, including detailed terms and conditions, more manuals, etc. or a salesman to call on you.

Sincerely,

Gordon Bell  
Vice President Research and Development

1 March 1987

nnn  
iiii  
sss  
tttt

Dear Computer Science Department head:

I built the first mini- and time-sharing computers including VAX as head of R and D at Digital Equipment. As the first Assistant Director for Computer and Information Science and Engineering at NSF, the most important problem for computing was closing the chasm between scientists and engineers who are entering the new era of supercomputing, and computers scientists who have the ability to lead in this revolutionary era.

Now, I would like to introduce you to Titan, the first, interactive Graphics Supercomputer.

Titan was designed specifically to revolutize science and engineering by providing exceptional power (10 Megflops on 100 x 100 Linpack - or roughly equal to a Cray 1 or IBM 3090), a large virtual and physical memory, and at less than \$100,000. It is designed to be used directly on projects. Furthermore it provides a "free" supercomputer's worth of power for graphics and visualization. It's a networked Unix machine with Cray and VMS Fortran, and C that's compatible with the Unix workstation environment.

While Titan may be the most important computer for scientists and engineers since the VAX and the Cray 1, it is equally important for computer science departments.

1. Titan is the ideal tool to provide a supercomputer for computational science training and research aimed at vector processing. Computer science must be restructured to reflect the major change in computer architecture with the introduction of vectors in supercomputers, mini-supercomputers, and now our graphic supercomputer.
2. Titan is an ideal system to introduce and explore the subtleties of parallel processing since it provides 1 to 4 vector processors. Primitives were added to Unix and the languages.
3. Titan renders 3D images of all types at the rate of 50 Million pixels per second. Furthermore we license Dore' (Dynamic Object-oriented Rendering Environment) to operate on all workstations.
4. By finally having supercomputer capacity, you can explore data-structures for real world objects including solids, atoms and molecules, fluids, visual images, animated figures, and mathematical functions.
5. It's completely affordable.

In summary, I feel we've built a tool that will revolution science and engineering, provided that computer science helps. Let me urge you to order one. If you don't like it we'll refund your money . A price list, using the NSF discount schedule is included.

Alternatively let me know if you want more information, including detailed terms and conditions, more manuals, etc. or a salesman to call you.

Sincerely,

Gordon Bell  
Vice President Research and Development



CSE

(over)

**NATIONAL SCIENCE FOUNDATION  
Computer and Information Science and Engineering  
1800 G Street, NW  
Washington, DC 20550**

January 30, 1987

Dear Colleague:

This letter is being sent to you in order to inform the scientific community of important activities presently taking place at the National Science Foundation. The fiscal 1987 Budget for the National Science Foundation includes funds of several million dollars in support of an initiative known as: "Computational Science and Engineering" (CSE). These funds are being distributed among the various disciplines: Biological, Behavioral and Social Sciences; Computer and Information Science and Engineering; Mathematical and Physical Sciences; Science and Engineering Education; Engineering; and Geosciences. It is anticipated that this new program will stimulate activity at the interface between the sciences and advanced computer technology. The NSF strongly urges investigators to inquire further about the details of the initiative with the various program directors at the Foundation. Enclosed with this letter is a program announcement (NSF 86-91) that describes the goals of the overall NSF/CSE programs.

Many of you may know that there have been a number of organizational changes at NSF. One is the creation of a new Directorate for Computer and Information Science & Engineering (CISE), which combines several preexisting computer activities from other directorates, the Division of Computer Research, the Division of Information Science and Technology, and programs in Computer Engineering, Communications and Signal Processing, and the Office of Advanced Scientific Computing.

CISE supports research in computer science, information systems and processing, robotics, networking and communications, microelectronics, advanced scientific computing and intelligent systems. The overall goal of the effort is to improve the knowledge base, research infrastructure and professional labor force needed to understand and improve the nature, synthesis and use of computing and information processing devices and systems. The current structure of CISE includes 5 divisions:

- Computer & Computation Research
- Information, Robotics & Intelligent Systems
- Advanced Scientific Computing
- Networking & Communications Research & Infrastructure
- Microelectronic Information Processing Systems

**FORMATION OF RESEARCH TEAMS**

Although many of the efforts described below can be performed by single investigators, and will be, to some extent, supported in that form, this new initiative will emphasize strong inter-disciplinary approaches to the enhanced computing capability and environment of the scientist and engineer. Proposals involving computer scientists, mathematicians, scientists and engineers, and specialists in such areas as computer graphics, might be integrated in such a way as to form an inter-disciplinary group or team, addressing specific problems of importance to one or more scientific or engineering disciplines. For example, such proposals might be strongly coupled with the efforts of innovators of state-of-the-art algorithms and software for application on machines with highly parallel architecture. Such approaches could develop new paths for entire disciplines to follow. They will be coordinated among CISE programs and the NSF scientific and engineering disciplines.

**RESEARCH OPPORTUNITIES IN CISE**

Proposals with a strong interdisciplinary approach are being encouraged in the following computational areas, although this list is not intended to be complete:

- Software and Algorithm Development
- Visualization, Graphics and Image Processing
- Network and Communication Systems
- Application of Advanced Technologies to problem solving
- Formation of Novel Computational Strategies
- Performance Evaluation of Computer Systems and Software
- Distributed and Parallel Processing and Vectorization

**Software and Algorithm Development:** It is widely accepted that software lags hardware development. This fact is especially true for supercomputers and other machines of advanced architectures. Within CISE, the Computational Science and Engineering Initiative will focus on research that addresses the development of novel algorithms and their implementation into useful software packages, the creation of friendly working environments, and the automatic production of fast, efficient code for scientists and engineers working on advanced computers. Innovations in languages, user-friendly interfaces, software tools, etc. might address issues related to the speed up of code development and therefore the productivity of researchers. Methods which assist in the portability of code across a variety of advanced machines will also be considered.

**Visualization, Graphics and Image Processing:** More powerful visualization capability is being demanded to take advantage of the most powerful machines. Substantial insights are already being gained from graphics, which is the only way to understand many scientific phenomena. Among the many research topics in graphics and image processing are: extemporaneous, interactive steering of numerically intensive calculations; dynamic visualization of fields in higher dimensions; high bandwidth graphics, networks and protocols; massive data set handling and standards; vectorization and parallelized algorithms for visualization; workstation-driven remote use of supercomputers; standard graphics-oriented scientific programming environments.

**Performance Evaluation:** A recent NAS/NRC report on "An Agenda for Improved Evaluation of Supercomputer Performance" remarks on the severe lack of scientific foundation, regarding our ability to evaluate the performance of advanced computers. Investigations into the definition and techniques for performance evaluation of parallel or other computer systems are encouraged either as the principal subjects of proposals, or as components of other research projects in this initiative.

**Distributed and Parallel Processing and Vectorization:** The direction of advanced scientific computing is clearly headed toward parallelism to achieve increased capacity. Since the complexities of programming in parallel environments with optimally vectorized code place even more challenging demands on software and algorithm development, the Computational Science and Engineering Initiative will emphasize means to provide effective scientific computing in vector and highly parallel environments. For example, the initiative will consider methods for automatically parallelizing existing scientific codes or rewriting them for efficient use on machines of advanced architectures. Also, software tools for increasing productivity of the programming environment on parallel and distributed architectures will be encouraged especially, for vector and multiprocessor computers.

**Advanced Technologies:** The Science and Engineering Initiative welcomes proposals concerned with areas of technology that have a strong impact on the conduct of future computing. Examples include high capacity and/or high performance mass storage coupled with appropriate file and data base management systems, optical computing, neural networks, non-binary computing, or any such ideas that could influence the nature of advanced scientific computing. The CSE Initiative will cooperate with other programs on the potential application of advanced computing technologies and systems to scientific and engineering problems. Proposals of this type will be coordinated as appropriate both within and outside the Foundation.

**Formulation of Novel Computational Strategies:** New computer architectures, communications technologies, languages, and other software or hardware advances becoming available offer promise of greatly enhanced speed, flexibility, or cost-effectiveness in performing scientific and engineering research. However, the hope for significant increases in insight to discipline specific problems may demand a fundamental revision in the strategic approach taken toward solving problems to make effective use of these options. Investigations into alternate ways of formulating and computing important scientific and engineering problems are encouraged.

**Network and Communication Systems:** Recently increased accessibility of advanced computing resources opens possibilities for new, computationally-based, advances in the understanding - i.e., analysis and especially design/synthesis - of computer networks and communication systems generally. This Initiative will entertain proposals for computational research in such problem areas as: event-based, Monte Carlo, or other simulation methodology applied to very large scale computer networks with attention to realistic detail; protocol design based on computational studies of state-machine models of networks with state spaces so large as to render such studies hitherto impracticable; specialized, interdisciplinary studies of Presentation- and Application-layer protocols; knowledge-based or other expert aids for intelligent dynamic network management; and research using symbolic computation in studies of algebraic coding theory. Proposals in these and other appropriate topical areas will emphasize the innovative computational nature of the proposed investigations, and may include the use of advanced (e.g., highly parallel) architectures in the research.

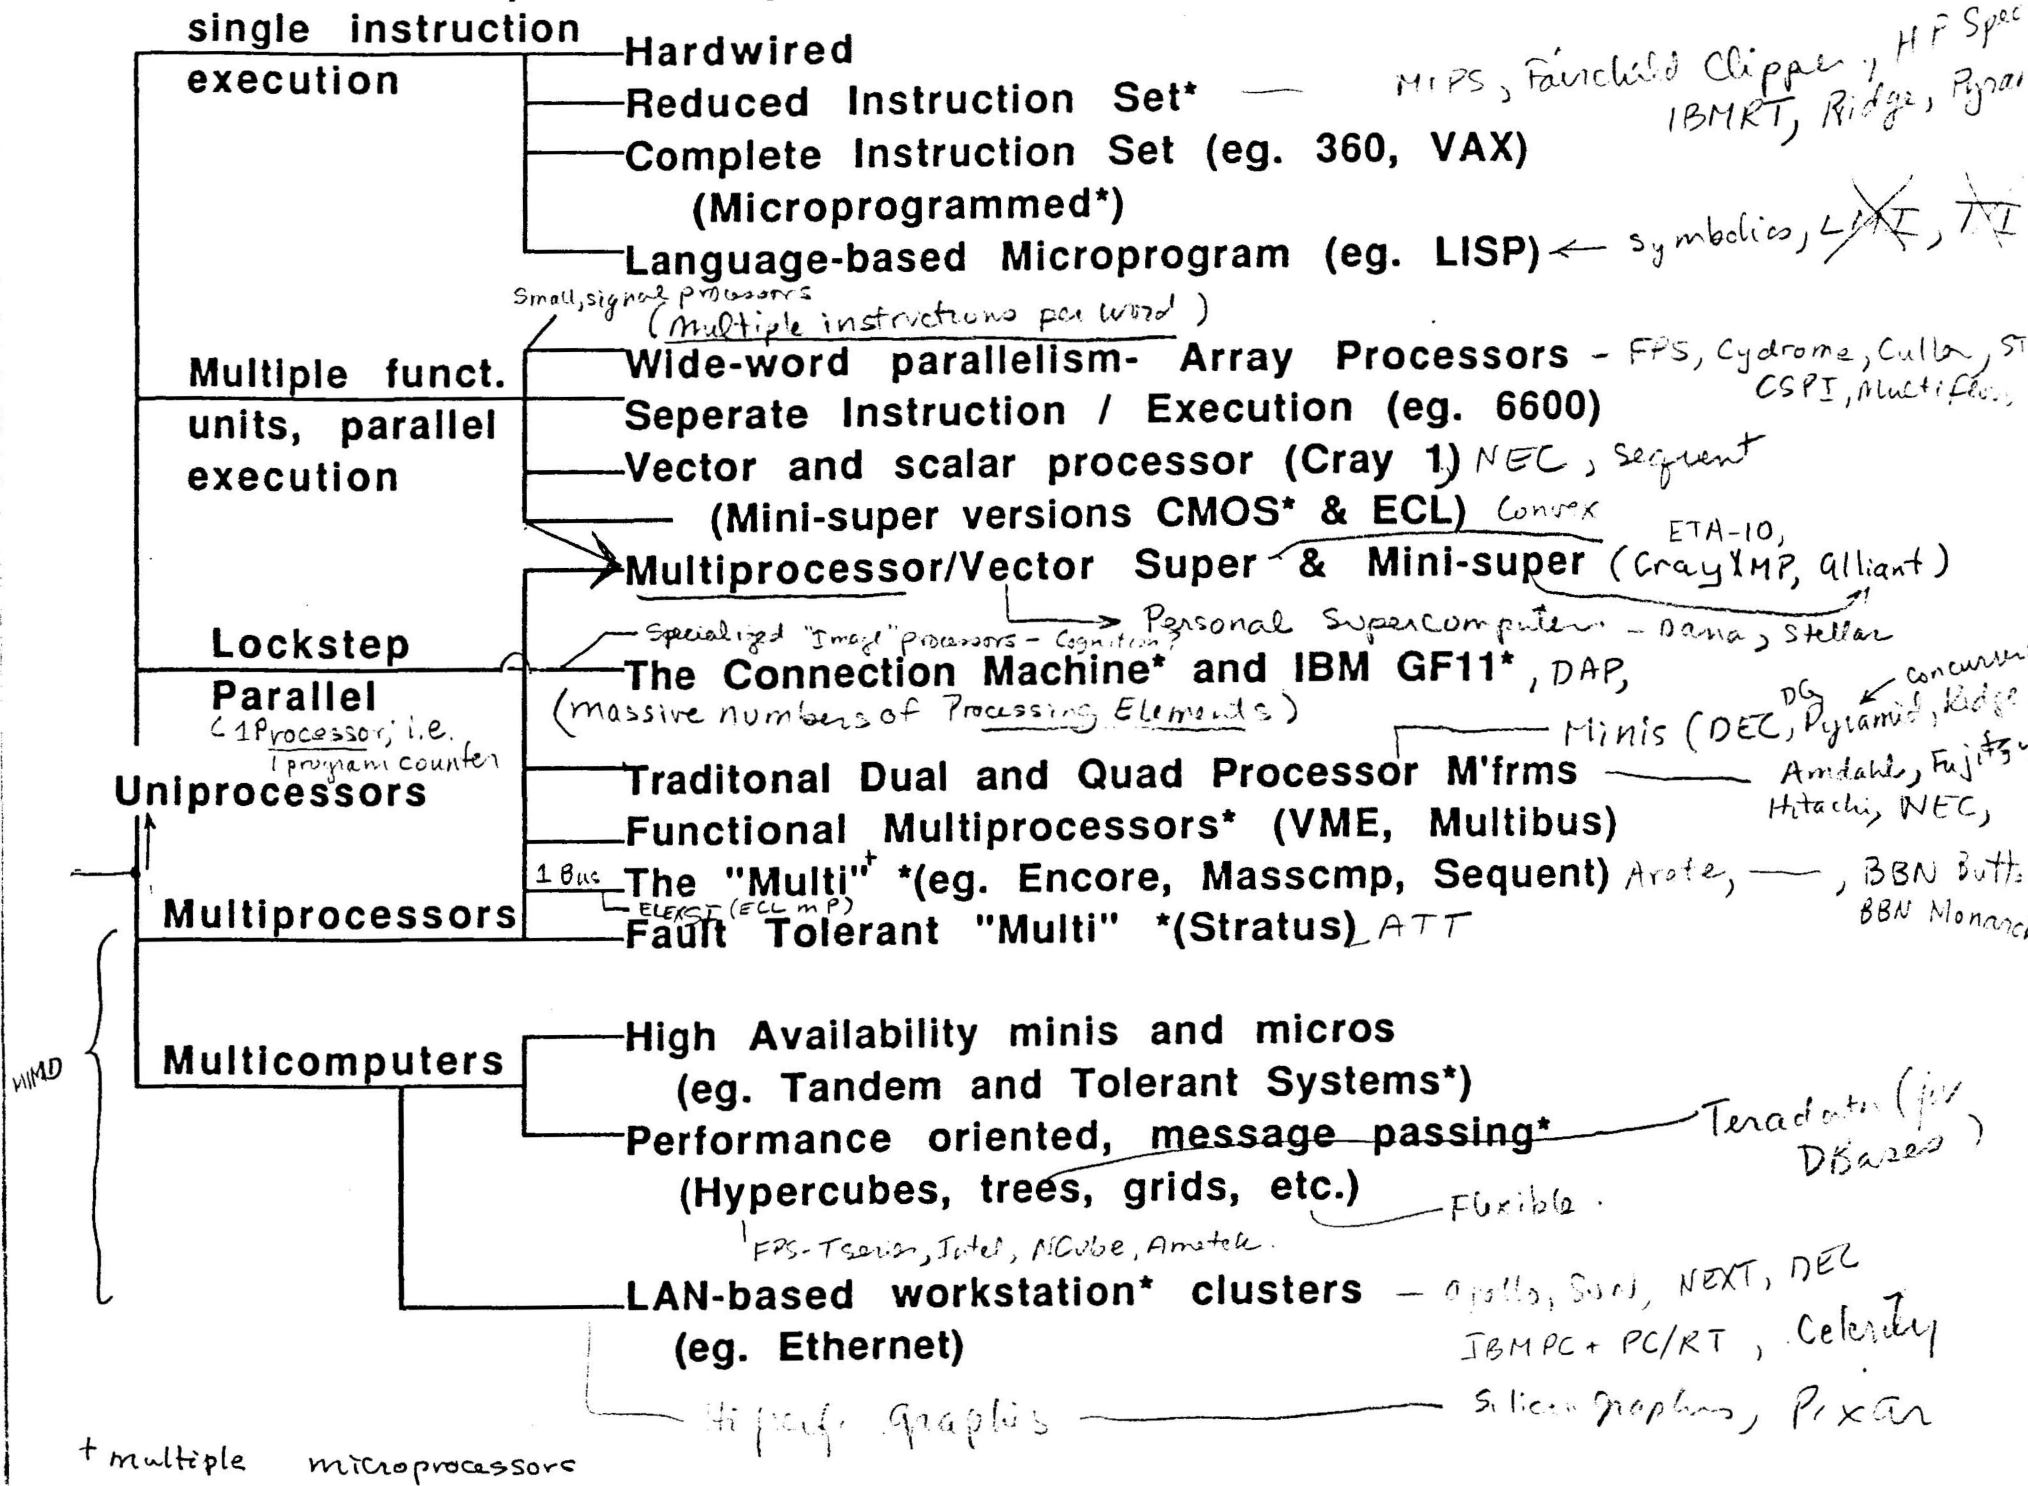
Sincerely,

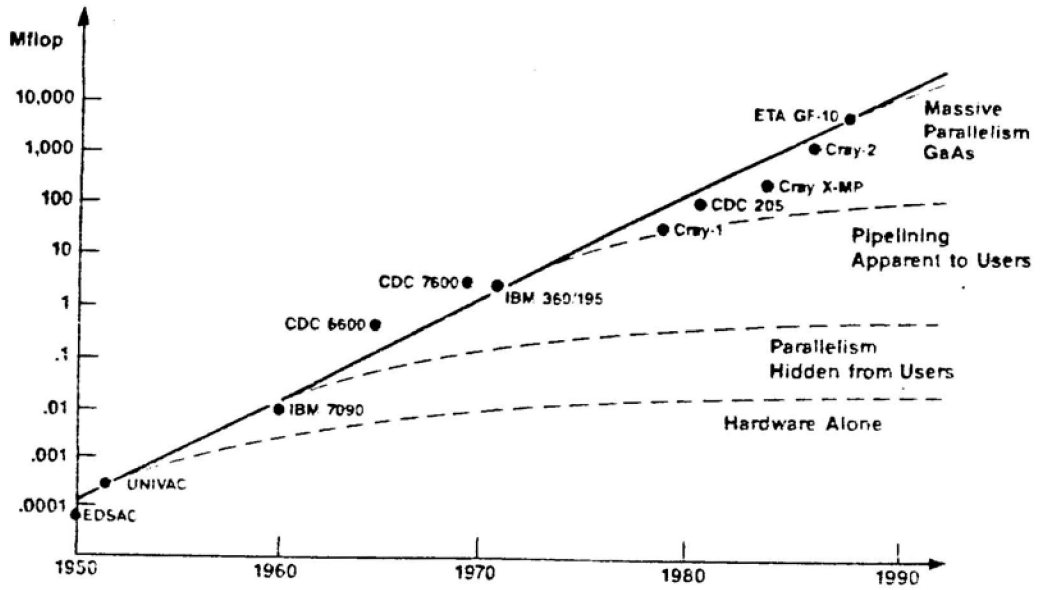


C. Gordon Bell, Assistant Director  
Computer and Information Science and Engineering




# Simplified Computer Taxonomy (Circa 1986)






## TRENDS IN THE EVOLUTION OF COMPUTERS



## PRICE/PERFORMANCE RATIOS AS FUNCTION OF SYSTEM SIZE

# PRICE/PERFORMANCE FOR DIFFERENT COMPUTER ARCHITECTURES






Fig. 1. Execution bandwidth of high-performance computers over the past 25 years.

Computers incorporating arrays of processors are beginning to emerge. Like vector computers, arrays of processors perform best with algorithms that have a minimal requirement for scalar computation. In addition, arrays of processors have limited data communication abilities between processors. Thus, when attempting to use them, one must seek algorithms that are highly parallel and with communication requirements that fit the architecture. The optimum arithmetic performance of these devices is potentially much greater than that of current supercomputers.

Because of the diminishing growth rate of performance in scalar processor architectures, we expect that the dominant architectural styles for future high-performance computers will include parallel operation or specialized pipeline structures. In both cases the programming tools and algorithms available to the user are inadequate and providing them is no mean task. It took between two and three years to adapt the programs developed for scalar

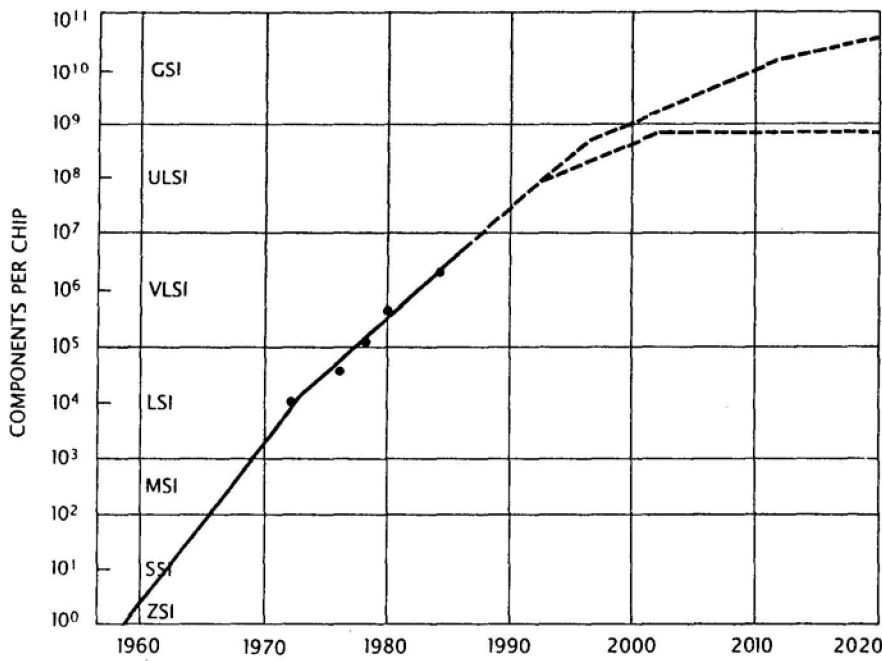
potential breakdown shakes so loose from vizir va- listi li- mc er

several ma- arsenide silicon in me small- tage may al factors. to switch ower to it, e buildup extremely switching city of the way from has three city of gal- ilicon de- h just as ostensibly

ierarchy balanced ere is this an at the s are nu- mpass all as addi- x device's e of n fi dentined to deter- e dimen- ion field-

consists oed to in- negative- ion) in a doped to of posi- ng these and the ion diox- ator. De- tween the etal elec- "input" electrons ion sub-

creasing density, as six chips made by the Fairchild m 1959 to 1985 show. At the top left (1959) is the first it (1961) is the first planar integrated circuit on a single along with other components. At the middle left (1964) linear r- ated circuit; it has five transistors. At the logi it has 180 transistors. At the bottom left inclu- entire central processing unit; it has 20,000 : (1985) is the CLIPPER CPU; it has 132,000 transistors.



**NUMBER OF COMPONENTS PER CHIP** doubled annually in the 1960's. In about 1972 designers ran out of unused space on the chip for additional components, and the rate fell somewhat. Nevertheless, according to the more optimistic projection, gigascale integration (GSI)—a one-billion-component chip—will be achieved by the year 2000. (The projections differ in assumptions about limits imposed by chip-fabrication processes.)


strate and the insulator. The electrons form an induced channel, allowing "output" current to flow from source to drain and sending a logic signal to the next stage of the circuit. In the absence of an input signal no channel is formed, and no output current results.

Since metal-oxide-silicon technology has become the dominant one throughout digital electronics, achieving GSI rests on the continuing attempt to scale down the MOSFET. The process of scaling down begins with the definition of a scaling factor, often called  $S$ . All the lateral and vertical dimensions of the MOSFET are then reduced by a factor of  $S$ . (Thus if the scaling factor were 2, the height and width of the device would be decreased to half their original values.) In addition, the supply voltage is reduced by the same factor, which keeps the strength of the electric

field constant and prevents an increase of the stress on the device.

The parallel reduction of size and electric-field strength yields some remarkable advantages. The time required to switch the device, which depends on the length of the channel, decreases by a factor of  $S$ . The power dissipated per unit of chip area remains constant, and so the problem of heat removal is not made worse. The packing density of the transistors on the chip, which depends on the area of the device, increases by a factor of  $S^2$ . Perhaps best of all, the energy consumed in each switching operation, which depends on the power and the switching time, decreases by a factor of  $S^3$ . Thus the result of scaling is a chip that has more devices switching faster and using less power to do so.

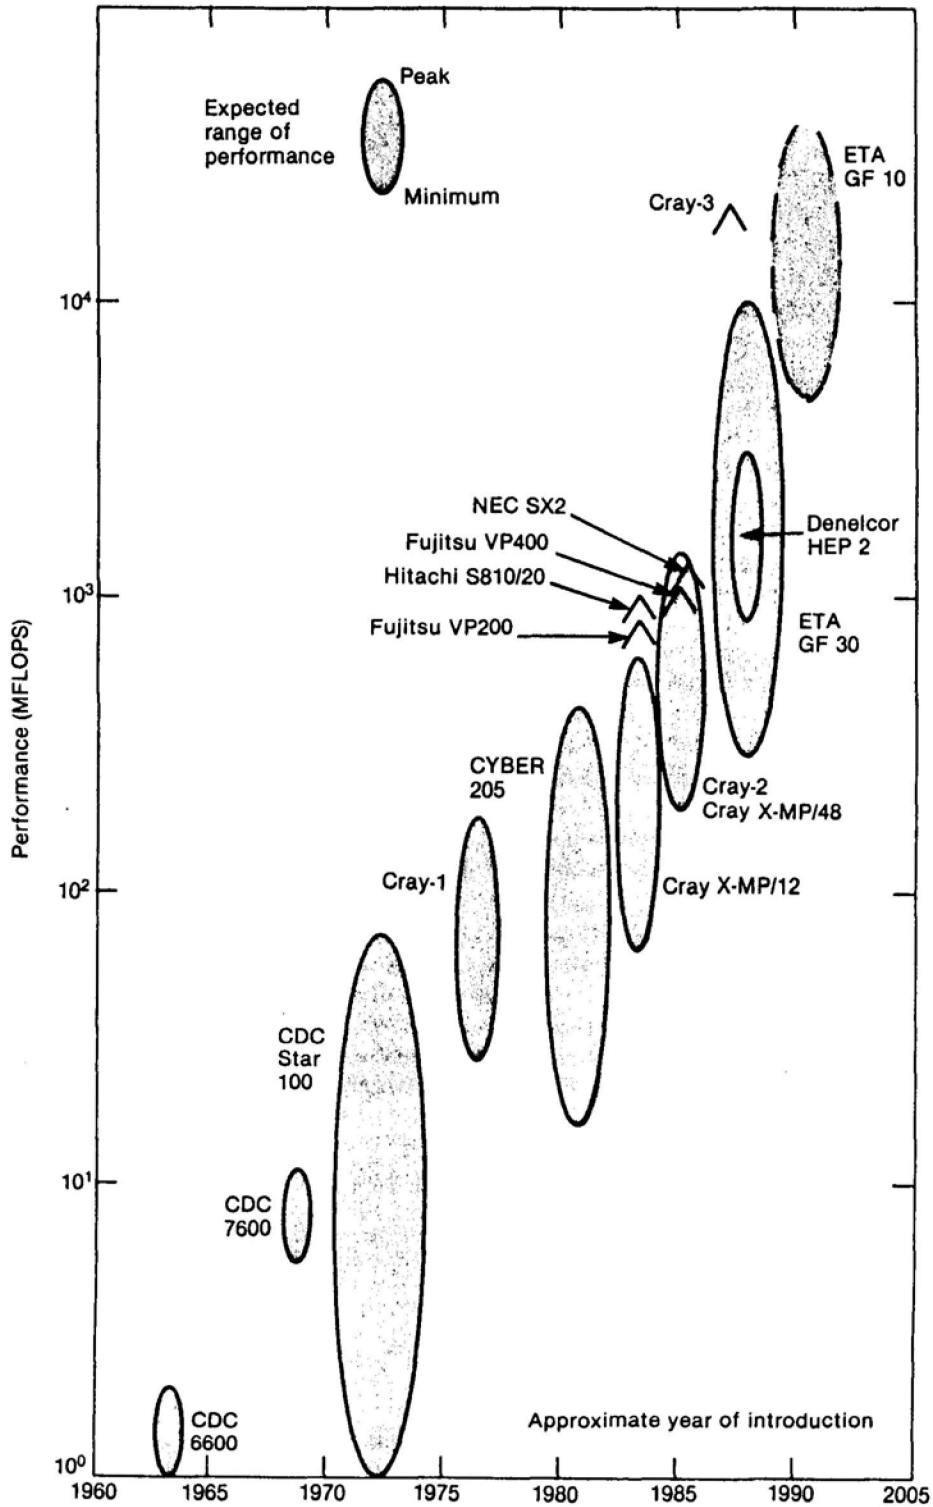
Given such advantages, the designer of chips would like to know how far scaling can go. The answer lies in the minimum allowable length of the MOSFET's channel. That minimum emerges in part from the interplay of the supply voltage and the doping concentration, which is the concentration of impurity atoms (usually boron) in the  $P$ -type substrate of the channel. At every junction between  $N$ -type and  $P$ -type materials, charge carriers migrate to the opposite side of the interface, where their concen-



EVOLUTION OF GENERAL-PURPOSE COMPUTING during a 40-year period is charted by colored bands. Four kinds of machines are tracked: mainframes (blue), minicomputers (red), personal computers (green) and embedded computers (yellow). Each of the bands defines the range of computing power, in millions of instructions per second (MIPS), that is available from a specific kind of machine at a particular time. The dotted lines represent projections beyond 1987. In any year the computing power of


mainframes is greater than that of minicomputers; the latter are more powerful than personal computers, which outperform embedded devices. Furthermore, computing power is cheaper on less powerful machines. For instance, in 1987 the approximate relative cost of executing one million instructions per second on a mainframe computer is 100 units; on a minicomputer it is 40 units, on a personal computer it is three units and on an embedded computer (if one were powerful enough) it would be .15 unit.

Figure 1.—Current and Projected Supercomputers, 1960-90



SOURCE: Sidney Fernbach.

**Figure 2.—Range of Federal Policies Possible With a Supercomputer 200 Times the Current Capabilities**




KEY:  
 [White box] = New science  
 [Light gray box] = More science  
 [Dark gray box] = Wider access/educational use



# New Languages for Old

by Arthur Fisher

The emerging shapes of scientific computers will stress existing languages and algorithms, perhaps beyond endurance.






Figure 1 Past Growth and Future Projections for Computer Speed and Computer Cost.




Figure 1 Development of Supercomputers - Computer Speed.  
(Dean Chapman)




Figure 2 Development of Supercomputers - Memory Size.  
(Dean Chapman)




Figure 11 Performance of CDC and CRAY High-end Computers Actual and Anticipated Through 1990. (For 60/64 bit arithmetic)




Figure 13 Cost Increase for Supercomputers.

NSF Center for CAI  
 Access

Phase I - 2 22 2 2 112

1/87

Sci/Eng Computers in U. S. Universities 84, 87 90

variable of ... in Cray 1's

D. Bell  
 1/87

	Sci, Pwr,	U84*	NSF**	U87	NSF	U90	NSF		U84	NSF	U87	NSF	U90	NSF	
Cray 1/S,M	1	1	2	1	1				1	2	1	1	0	0	
Cray 48 & 2	8	0	0	1.75	4	3	1		0	0	14	32	24	8	
Cray YMP	32						3		0	0	0	0	0	96	
CDC 205	1.4	3	0	5	1	5	0		4	0	7	1	7	0	
ETA 10	20					2	1		0	0	0	0	40	20	
IBM 3090/200	2			40	2	150	4		0	0	80	4	300	8	? for sci.
Total Supers		4	2	47.8	9.8	160	9		5	2	102	38	371	132	
IBM 3080	0.17			80		20			0	0	14	0	3.4	0	? for sci.
Alliant	0.67			12	1				0	0	8	1	0	0	
Convex	0.5			12					0	0	6	0	0	0	
SCS	0.67			2	1				0	0	1.3	1	0	0	
Tot. Mini-sup		0	0	26	2	0	0		0	0	15	1	0	0	
DEC 7xx	0.011	1450		2650					16	0	29	0	0	0	
DEC 86xx	0.075	3		280	4				0	0	21	0	0	0	
IBM 4381-13	0.1			250					0	0	25	0	0	0	
DEC MicroVAX	0.01			1600					0	0	16	0	0	0	
SUN 3/260	0.038								0	0	0	0	0	0	
Total W/S's															
FPS 164	0.14			19	3				0	0	2.7	0	0	0	
FPS 164Max	1			10					0	0	10	0	0	0	
FPS 264	0.5			12	5				0	0	6	3	0	0	

\* university-provided ... NSF centers and ... NCAR

**Worlton and Associates**  
3089 Villa, Los Alamos, NM 87544  
(505) 662-4011 or 662-2724

February 26, 1987

Dr. James Decker  
ER-2, DOE  
1000 Independence Avenue, SW  
Washington, DC 20585

Dear Jim:

I am transmitting herewith the report providing an update on Japanese/USA supercomputer competition that you requested for the FCCSET committee. It pains me to send it out in this rough draft form, but the February 28th deadline sets the ultimate limit to further work. I will send a copy to John Riganatti, who will also be sending you his report. John and I talked about merging our efforts into a more complete report on this subject; perhaps FCCSET would like a more complete and polished report. If so, let us know.

Sincerely yours,

  
Jack Worlton

wjw

Encl: a/s

cc: J. Riganatti, SRC  
R. Ewald, CRI  
T. Vacca, ETA  
A. Lundy, LANL, IT-3  
N. Morse, LANL, C-DO  
File: Japan

RECEIVED  
OASD/CIS  
Date MAR 10 1987



ROUGH DRAFT FOR COMMENT  
2/26/87

**AN UPDATE ON  
COMPETITION IN THE SUPERCOMPUTER INDUSTRY:  
JAPAN vs USA**

by  
**Jack Worlton**  
Los Alamos National Laboratory

**Submitted to:**  
**The Federal Coordinating Council**  
**on**  
**Science, Engineering, and Technology**  
**Washington, D.C.**  
**February 28, 1987**

## TABLE OF CONTENTS

1. EXECUTIVE SUMMARY .....	1
2. INTRODUCTION .....	4
<u>2.1 Purpose</u> .....	4
<u>2.2 Scope</u> .....	5
<u>2.3 Method</u> .....	5
<u>2.4 Taxonomies</u> .....	5
3. MARKET ISSUES .....	7
<u>3.1 Market share</u> .....	7
<u>3.2 Timing</u> .....	7
<u>3.3 Compatibility</u> .....	8
<u>3.4 Market strategies</u> .....	8
<u>3.5 Architectural issues</u> .....	9
<u>3.6 Product comparisons</u> .....	9
<u>3.6 Summary</u> .....	10
4. TECHNICAL ISSUES .....	12
<u>4.1 Component technology</u> .....	12
<u>4.2 Architecture</u> .....	12
5. STRUCTURAL ISSUES .....	14
<u>5.1 The Japan Problem</u> .....	14
<u>5.2 Government relations</u> .....	15
<u>5.3 Cultural driving forces</u> .....	16
6. SUMMARY AND CONCLUSIONS .....	17
7. REFERENCES AND BIBLIOGRAPHY .....	18

## 1. EXECUTIVE SUMMARY

Purpose, scope, and method. This report is being prepared at the request of Dr. James Decker, on behalf of the Federal Coordinating Council on Science, Engineering, and Technology (FCCSET). The deadline for the report is February 28, 1987, a month in which the author is on travel for 20 of the 28 days, so this report is necessarily brief and incomplete (including incomplete editing). The method employed in this report is that of a situation audit, which is conceptually a matrix: on the left side are strengths plus resulting opportunities and weaknesses plus resulting problems; on the top are the relevant firms from the U.S. (Cray Research, Inc. (CRI), and CDC/ETA Systems) and Japan (Fujitsu, Hitachi, and NEC). We have not included IBM in this report, because their position is somewhat ambiguous: the IBM 3090VF is not considered a supercomputer by IBM for purposes of avoiding export controls, but it is considered a supercomputer by IBM when certain customers are being approached. In the long term, IBM should probably be included in this kind of analysis.

Background. In 1983 two Japanese firms (Fujitsu and Hitachi) began early deliveries of supercomputers in Japan, and in 1985 NEC began deliveries of their supercomputer. Prior to this time, the supercomputer market had been a small and exclusively American market, with only Cray Research and CDC offering such computers. In 1983 CDC formed a new small company, ETA, to develop their next generation of supercomputers; CDC is the primary stockholder in ETA, so this company is sometimes referred to as CDC/ETA or simply as ETA. The American and Japanese firms are quite different: CRI has revenues of a few hundred million dollars per year and ETA is still just a startup company, whereas the Japanese firms have annual revenues on the order of \$10 to \$20 billion. The American firms are not semiconductor manufacturers and must depend on other companies for their components, whereas the Japanese firms are all world leaders in merchant semiconductor manufacturing. On the other hand, because the American firms were in this market before the Japanese firms, they had a market share advantage, so the problem for the American firms was to hold their market share, whereas the problem for the Japanese was a matter of penetrating an existing market. The Japanese firms enjoy several structural advantages, the most important of which is the "partnership economy" of Japan, in which the government and industries of Japan work as partners in promoting industrial expansion. In the U.S., the government and industry are often adversaries, with the most obvious example being export controls which constitute a major problem for the American firms but not for the Japanese firms. Since their entry into this market in 1983, the Japanese firms have been most successful in Japan, with Fujitsu leading in installing computers there, at last report about 30. None of the Japanese firms has so far been particularly successful in the international market; Fujitsu has been the most aggressive, but NEC (whose motto is *Attack!*) is now supplanting Fujitsu as the most aggressive

in international supercomputer marketing. Hitachi has limited their supercomputer marketing to Japan.

Situation audit. The reasons for the ability of the American firms to hold their market share in the face of competition from such strong Japanese firms include (1) market share, (2) technical leadership, and (3) an existing marketing infrastructure. All other things being equal, a large market share tends to be self-perpetuating because of repeat orders, compatibility considerations, the desire of customers to collaborate with other sites having the same type of computers, and a large and stable base of system and application software. The task facing the Japanese firms is to attack the phrase "all other things being equal." The strengths of the Japanese firms that might change this assumption include their strong semiconductor development capabilities, which can be used to develop advanced components one to two years before the American firms, and their financial strengths, which can be used to "buy" contracts away from the smaller American firms who cannot afford the heavy discounts being offered by the Japanese firms. If this competition were being conducted wholly within the U.S., antitrust laws would prevent the Japanese firms from using many of their marketing strategies such as "anticipatory pricing" (selling below cost), but in the international market, no such rational protections exist.

Specifically, the strengths of Cray Research include a market share of about two-thirds of current supercomputer installations, technical leadership in parallel processing which they began in 1982, a rich base of application and system software, and a strong marketing and technical support infrastructure. Their primary weakness is their dependence on other firms to provide the high-performance logic and memory components they need for new generations of supercomputers. CDC/ETA has a smaller, but not inconsiderable market share of some thirty-odd machines, which will be an advantage to them in marketing their new machine, the ETA<sup>10</sup>. They have long experience in developing supercomputers and hence a strong and knowledgeable staff. Their weaknesses include the fact that they have no current product to sell, with the ETA<sup>10</sup> hardware and software still being in development and their Cyber 205 being obsolescent; their status as a startup company with little or no income; and their dependence on other firms for advanced components.

The specific strengths of the Japanese firms include their leadership in semiconductor manufacturing, which should give them a timing advantage in developing new generations of supercomputers; the support they receive from the Government of Japan in the form of government-supported research projects and avoidance of export licensing problems; and their financial strengths which can be used to support deep cuts in pricing (referred to as "anticipatory pricing"). Their weaknesses include a small market share, immature software, incompatibility with most supercomputer users, lack of credibility for software and maintenance support, and architectural obsolescence (marketing serial processors in a world rapidly moving toward parallel processing).

Summary. The "first round" of international competition in supercomputing must be conceded to the American firms, primarily Cray Research. Japanese successes have been limited largely to Japan where cultural preferences have made it easy for the Japanese firms to corner that market. There are widely circulated rumors of a new generation of supercomputers from the Japanese firms being introduced in 1987, although no formal announcements have been made. Both Cray Research and ETA Systems also plan to offer new computers in 1987, so it is possible the "second round" will still be a wash, unless there are some big surprises from the Japanese firms. It is in the long term that the Japanese advantages of component development, financial strengths, and government relations will be most evident. To survive, the American firms must somehow gain access to timely development of high-performance logic and memory components independent of Japanese firms, and it is not yet clear how they will do that. Collaboration with some small "niche vendors" seems to offer the best hope at the moment. The American firms also need to somehow counter the Japanese advantage in government relations. Whether this should take the form of direct government support as in Japan, or merely removing governmental barriers such as export control delays, is a topic on which there is agreement only on the latter point.

## 2. INTRODUCTION

2.1 Purpose. The purpose of this report is to conduct a brief "situation audit" of competition between Japan and the USA in the supercomputer industry. Prior to 1983, the supercomputer market had been exclusively American for about twenty years, with the last supercomputer marketed by a foreign country being the British ICL Atlas from the early 1960s. However, in 1983, two Japanese companies, Fujitsu and Hitachi, began deliveries in Japan of supercomputers whose performances were within about a factor of 2 of the performance of the leading American supercomputers, and in 1985, NEC began deliveries of computers that were faster than the Fujitsu and Hitachi machines by about a factor of 2. As shown in Table 1, these products--the so-called "Developing Generation"--were the result of earlier developments--the so-called "Embryonic Generation"--by these companies."

	EMBRYONIC GENERATION	DEVELOPING GENERATION
FUJITSU	<ul style="list-style-type: none"> <li>• FACOM 230-75 APU</li> <li>-- 1977</li> <li>-- 22 MFLOPS</li> <li> </li> <li>-- Shared main memory</li> <li>-- Fully pipelined</li> <li>-- 1972 vector registers</li> <li>-- AP Fortran</li> </ul>	<ul style="list-style-type: none"> <li>• VP-100/200/400</li> <li>-- 1983</li> <li>-- 250/500/1000 MFLOPS (peak)</li> </ul>
HITACHI	<ul style="list-style-type: none"> <li>• M-180 and M-200 w/IAP</li> <li>-- 1978 and 1980</li> <li>-- 24 MFLOPS</li> <li>-- Vectorizing compiler</li> <li> </li> <li>-- Shares func. units w/CPU</li> </ul>	<ul style="list-style-type: none"> <li>• S810/10 and 20</li> <li>-- 1983</li> <li>-- 315/630 MFLOPS (peak)</li> </ul>
NEC	<ul style="list-style-type: none"> <li>• ACOS 1000 W/IAP</li> <li>-- 1981</li> <li>-- ?? MFLOPS</li> </ul>	<ul style="list-style-type: none"> <li>• SX-1 and SX/2</li> <li>-- 1983</li> <li>-- 570/1300 MFLOPS (peak)</li> </ul>

Table 1. Embryonic and Developing Generations of Japanese Supercomputers.

The "IAP" refers to an "Integrated Array Processor" that was an arithmetic accelerator attached to the mainframes. Thus, these companies did not, as often believed, suddenly begin producing supercomputers, but had been working on vectorizing units and their software for several years. Both of these generations borrowed ideas from prior American designs, including



the CDC Star-100, the Cray-1 and the CDC Cyber 205. Since the introductions of the Japanese products in 1983 and 1985, several products have been added on the low end of the cost and performance range.

It is anticipated that the third generation of Japanese supercomputers--presumably a "Mature Generation"--will be forthcoming in the next year or so. Very little information but lots of rumors have been circulated concerning these machines.

2.2 Scope. The scope of this report is severely limited by the time constraints allowed for its preparation, so it is mostly an incomplete digest rather than a complete and detailed report. Further information can be obtained by checking the sources noted in Section 7, References and Bibliography.

2.3 Method. This report is a situation audit, with emphasis on key issues facing the competitors and their host countries. We include not only the usual marketing and technical issues, but also the "structural" issues that are crucial for understanding any competitive situation with respect to Japan. This is often referred to as "The Japan Problem."

2.4 Taxonomies. To clarify the class of computers under discussion, we include two taxonomies. Figure 1 is a partial taxonomy of high-performance scientific computers that shows the three main categories of such machines: research, special-purpose, and general-purpose. We shall be concerned here with general-purpose high-performance computers. Within that category are three types of computers: supercomputers, high-end mainframes, and "mini-supers." Although there is some overlap in the performance ranges of these types of computers, the supercomputers as a class outperform the other two types of high-performance computers and this category is the subject of this report.

Within the category of supercomputers, there are three "classes" often referred to, as shown in Figure 2. The performance ranges shown are only approximate, of course. Supercomputers by their nature have very broad performance ranges compared to other kinds of computers, and the overlap of the performance ranges for the three classes is deliberate. Whereas the first two classes of supercomputers represent machines that have already been delivered (or are reasonably close), none of the Class 7 machines have been delivered, and these are merely announced plans of the companies whose products are shown.. The new generation of Japanese supercomputers will presumably fit into the Class 7 category.

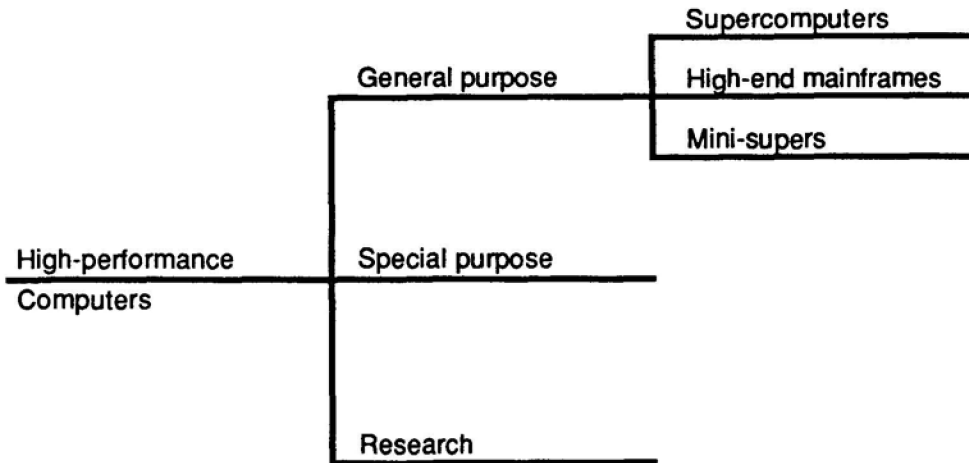


Figure 1. Partial Taxonomy of High-Performance Computers.

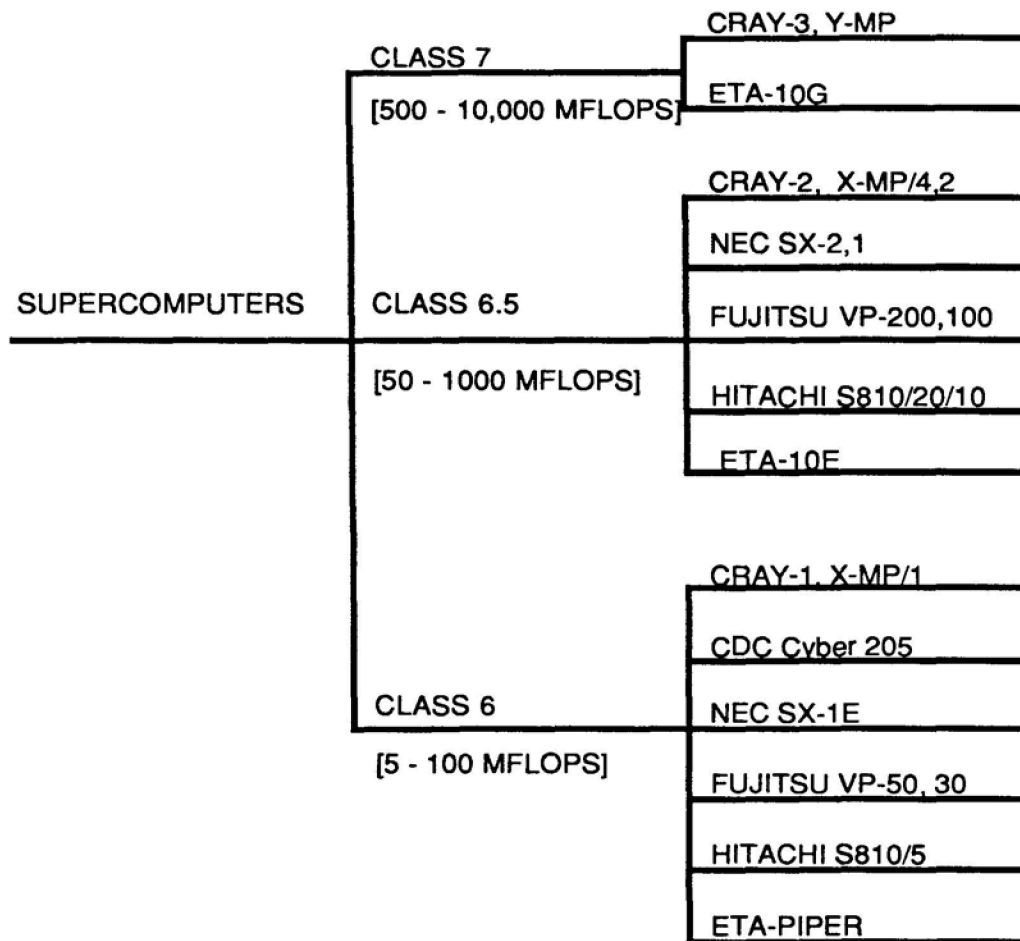


Figure 2. Taxonomy of supercomputers.



### 3. MARKETING ISSUES

3.1 Market share. Figure 3 shows the share of the supercomputer market held by Cray Research, CDC/ETA, and the Japanese companies as of mid-1986.

#### SUPERCOMPUTER MARKET



Figure 3. Shares of the supercomputer market as of mid-1986.

The percentages shown here are continually changing, of course, but it is roughly true that Cray Research sells about two-thirds of the supercomputers in the world, and the other third is divided between the three Japanese companies and CDC/ETA. Or, to put it another way, since the Japanese began delivering supercomputers in 1983, they have captured roughly 18 percent of the market. However, a basic principle of marketing is that *a large market share tends to be self-perpetuating*, because of the commitment of the customer to the particular product, where in this context "commitment" includes such things as applications codes, user competence and training, operational skills, and site installation. These commitments are reflected in repeat orders, a rich body of application and system software, and collaboration with other sites using compatible computers. And while this is true of computers in general it is true of supercomputers in particular. The reasons for this lie in the effort needed to prepare applications software for these computers and the rich set of software that the customers can obtain with a minimum effort or expense. More software exists for the computers having large market shares because there are more users and more third-party software vendors developing such software for these computers. This is true of IBM's large share of the mainframe market and Digital Equipment Corporation's share of the minicomputer market, as well as Cray Research's share of the supercomputer market. Thus, Cray Research's large market share is one reason the Japanese companies have not made more progress than they have in penetrating the supercomputer market.

3.2 Timing. Another reason Cray Research has been able to hold off this intense competition from powerful Japanese companies is found in the timing of recent introductions. Cray Research

introduced the Cray-1 in 1976, and this machine had essentially no competition for about five years and only minimal competition for some after that until the internal competition created by the introduction of the Cray X-MP/2 in 1983. The Japanese vendors targeted the Cray-1 as the computer their computers should exceed in performance. However, by the time Fujitsu and Hitachi entered this market in 1983, Cray Research had introduced in 1982 a newer and more powerful product line, the Cray X-MP/1 and Cray X-MP/2, with one and two processors, respectively. The single-processor X-MP has about the same performance as the Fujitsu VP-200 in general-purpose computing, and about twice the performance of the Hitachi S810/20. Thus, these Japanese computers were indeed faster than the Cray-1, but this was no longer the relevant comparand by the time the Japanese products were introduced.

A similar situation occurred when NEC introduced the SX-2 in 1985. This computer was faster than the other Japanese supercomputers by about a factor of 2. However, by the time NEC introduced this computer Cray Research had introduced both the four-processor X-MP/4 and the four-processor Cray-2, thereby effectively preempting the NEC introduction.

During this period, Control Data Corporation, the vendor of the Cyber 205 (a computer roughly in the same performance category as the Cray-1) formed a new subsidiary, ETA, to design the next generation of supercomputers, generically referred to as the ETA<sup>10</sup>. These computers are in the late stages of their development, and ETA is expected to begin deliveries of hardware and software for the ETA<sup>10</sup> in the next twelve to eighteen months.

3.3 Compatibility. A commonly used guideline in the management of scientific computing is that *an incompatible computer must provide a performance gain commensurate with the cost of conversion*. This is usually quantified as a factor of 2, i.e., an incompatible computer must outperform a compatible computer by at least a factor of 2 to justify the cost of conversion. The Japanese computers failed to meet that criterion even with respect to a single-processor X-MP, let alone the dual-processor X-MP. Thus, it was not surprising that neither the Fujitsu nor the Hitachi products were able to penetrate this market except in Japan, where the well-known Japanese antipathy toward foreign products led to the acquisition of mostly Japanese supercomputers in spite of normal computer evaluation criteria.

3.4 Market strategies. The Japanese employ two distinctive strategies when attempting to penetrate a new market: targeting and anticipatory pricing. Targeting is a national industrial strategy and anticipatory pricing is a corporate strategy. Targeting refers to the practice of bringing overwhelming national resources to bear against a specific industry of another nation, such that the target industry is at a disadvantage. For example, the Japanese Ministry of International Trade and Industry's Super-Speed Computer System Project, for the period 1981 to 1989 and funded with \$100 million from the Japanese national budget, brought together the resources of Japan's six largest computer companies in a national project to develop the supercomputer technology that

would allow Japanese companies to become world leaders in this field. The American companies in the supercomputing industry, Cray Research and ETA, have revenues of less than 1/10 to 1/40 of the revenues of the leading Japanese companies, and the aggregate resources of the six major Japanese companies are even more overwhelming.

The financial strengths of the Japanese companies make it possible for them to employ the second of these strategies, "anticipatory pricing." Essentially this means that a company sells its products either below cost or at huge discounts, attempting thereby to attract customers away from the companies whose products are priced to make a normal profit; this is sometimes referred to as "dumping." The word "anticipatory" refers to the expectation that in the long term, as the Japanese company gains market share, their prices will be adjusted to generate profits. This strategy cannot be employed by small companies that must make profits in order to survive, but only by large companies with other divisions whose profits support this penetration of a new market. The anticipatory pricing strategy is being used currently by Japanese firms offering huge discounts, a case in point being the well-known sale of NEC's SX-2 to the Houston Area Research Council (HARC) [4,5]; similar efforts are occurring in other nations. Briefly, this strategy is an attempt on the part of a large company to "buy" a market away from a small company and thereby put the small company out of business. NEC's company motto of *Attack!* is well illustrated by this strategy.

3.5 Architectural issues. The Japanese supercomputers are all single-processor designs, and this has probably had some negative effect on their marketing efforts. There is a broad consensus among the world's computer scientists that computers of the future, and especially supercomputers, will be built using multiple processors, so acquiring one of the Japanese machines has meant a customer was buying "instant obsolescence" in the architectural sense. Not many customers want to spend the millions of dollars supercomputers cost without getting a current design.

3.6 Product comparisons. Table 2 lists current and projected supercomputer products as of February 1, 1987. A few of the "Next Generation" products may be shipped in 1987, but substantial customer shipments are not expected until 1988. This is also true of the "Future Generation" for the years 1988 and 1989. Table 3 shows some general characteristics of some representative supercomputers from the current generation .

VENDOR	CURRENT GENERATION (Early 1987)	IN DEVELOPMENT (1987-88)	FUTURE GENERATIONS (1989 or beyond)
Cray Research	Cray-1 Cray X-MP/1,2,4 Cray-2	Cray Y-MP/8 Cray-3	Cray MP
CDC/ETA	Cyber 205	ETA-10G ETA-10E ETA Piper	Unknown
Fujitsu	VP-30,50,100 VP-200,400	Unknown	Unknown
Hitachi	S810/5,10,20	Unknown	Unknown
NEC	SX-1E,1,2	Unknown	Unknown

Table 2. Supercomputer generations.

SYSTEM	FCS*	Cycle Time (ns)	No. PEs	Main Memory (MW)	Extended Memory (MW)
Cray-1	1976	12.5	1	1-4	----
Cray X-MP/1,2,4	1982,1984	9.5/8.5	1,2,4	1-16	32-512
Cray-2	1985	4.1	4	256	----
CDC Cyber 205	1981	20.0	1	4-16	----
Fujitsu VP-200	1983	14/7	1	8-32	----
Hitachi S810/20	1983	14	1	4-32	32-128
NEC SX-2	1985	6	1	16-32	16-256

\*FCS = First Customer Shipment

Table 3. General characteristics of some representative current-generation supercomputers.

**3.7 Summary of marketing issues.** In summary, Cray Research has been able to withstand the attacks of the larger Japanese companies during the past three years by virtue of its large market share, by its timely introduction of new products, by leadership in parallel processing, and by the incompatibility problem faced by the Japanese. ETA Systems is just now in the final phases of

product development, and it will about twelve to eighteen months before it is known how well they will do against their competitors, both domestic and foreign.

The continuing marketing problems facing the American competitors include (1) their limited ability to match the Japanese semiconductor-development capability, (2) the targeting strategy employed by the partnership between Japanese industry and the Government of Japan, and (3) the marketing strategies of anticipatory pricing and dumping employed by the much larger Japanese firms.

## 4. TECHNICAL ISSUES

### 4.1 Component technology.

The three Japanese supercomputer vendors, NEC, Hitachi, and Fujitsu, rank number 1, number 2, and number 7 in the world, respectively, as merchant semiconductor manufacturers [31]. The advantage this gives them in developing supercomputers is largely one of timing. They can develop new generations of advanced components about one to two years ahead of their American competitors, according to Tony Vacca, Vice President for Technology at ETA Systems. Suppose, for example, that a Japanese firm begins marketing a supercomputer with a 1 nanosecond (ns) cycle time two years ahead of a similar product from American firms. During this timing gap, there will be some market penetration before the American firms catch up, and after a few rounds of this experience, the total market would inevitably be captured by the Japanese. At the moment the fastest cycle times are found in the Cray-2 (4 ns) and the NEC SX-2 (6 ns). However, the Cray-2 issues instructions only every other cycle, so for scalar work its cycle time is more like 8 ns and only in long vectors does it appear as a 4 ns cycle. Thus, for practical purposes the Cray-2 cycle time is probably best thought of as about 6 ns. There are prospects for improving cycle times to 3 ns in silicon, 2 ns in gallium arsenide, and 1 ns in HEMT (high electron mobility transistors). For the American firms the problem is that the high-performance semiconductor market is so small that it attracts little attention from the major American semiconductor firms. Some small "niche" vendors who are willing to take a somewhat larger share of a small market have recently shown interest in serving this need, but they are evidently somewhat behind the Japanese firms in developing production versions of these advanced components.

### 4.2 Trends in supercomputer architecture.

There is an international consensus among computer scientists that the future of supercomputing lies with parallel processing, i.e., designs having multiple processors that can be used to shorten the solution time for a single problem. Cray Research began deliveries of such designs in 1982 with their Cray X-MP/2 with two processors, and in 1985 with their Cray X-MP/4 and Cray-2, each with four processors. The Cray Y-MP/8 will have a 5 ns or 6 ns clock and eight processors, planned for initial delivery in 1987; the Cray-3 is projected to have a 2 ns clock and 16 processors, for delivery in 1988 or 1989. ETA has parallel processors under development: the ETA<sup>10</sup>-G with a 7 ns clock and up to 8 processors; the ETA<sup>10</sup>-E with a 10.5 ns clock and up to 4 processors; and the ETA Piper with a 21 ns clock and 1 or 2 processors. There is, of course, the usual uncertainty about when systems under development will actually be delivered to customers in substantial quantities.

None of the Japanese firms has yet announced a parallel processor, although it is known that all of them are doing research on this kind of design, including the work being done on the Super-

Speed Computer System [28] sponsored by the Government of Japan. There have been some reports that this project is having problems, however [17]. In this sense, the Japanese supercomputers are architecturally obsolete, and this may have hindered some customers from taking an interest in these computers.



## 5. STRUCTURAL ISSUES

By "structural" issues we refer to those constraints on the supercomputer market that are built into the competitive environment over which the individual companies have little or no control, including the general trade relationships between Japan and the USA, cultural differences, and relationships between industry and government in the two nations.

### 5.1 The "Japan Problem."

An excellent summary of this problem is contained in a recent article by Karel G. van Wolferen, a Dutch writer who has lived in Japan since 1962 [1]. One of the myths about Japan is that it is a sovereign state like others among the Western nations. In other nations there is a source of power that can take responsibility for decisions and actions, but this is not true in Japan. Rather, there are *three* sources of power, none of which can assert that "the buck stops here," as did the American president Harry Truman. In Japan, the buck doesn't stop, it circulates among the politicians, the bureaucracy, and the industrialists. The evidence of this can be clearly seen in Nakasone's largely unsuccessful efforts to bring changes to the Japanese relationships with other nations. Western negotiators often express frustration at the seeming insincerity of Japanese negotiators because they say one thing during negotiations but do not follow through on apparent agreements when they get home. The problem here is caused not by Japanese insincerity but by a lack of understanding on the part of Westerners about the Japanese culture. Decision making on major issues can occur only by consensus among the three major centers of power, and thus it is impossible for any single party to represent the views of all three until later in time when a consensus has been reached. The point here is that Westerners should understand that Japanese politicians cannot make commitments in the same way politicians do in other Western governments.

A second myth about Japan is that its economy is market driven as are those of the other Western nations. Japan's economy falls into neither the free-market category nor the centrally-controlled category, but into what might be called a "partnership economy," where the partnership referred to is between the government and industry. As van Wolferen writes, "...it is impossible in Japan to separate the state from the socioeconomic system." And while it is true that the state is somewhat involved in the socioeconomic system in most nations, the involvement in Japan occurs to a degree that exceeds the involvement of any other nation, except for the centrally controlled economies. Unlimited industrial expansion is the consensus industrial policy of this partnership, to the single-minded exclusion of all else. In the United States and most other Western nations, other objectives such as defense and social-agenda items compete with industrial expansion, but not in Japan. Roadblocks to industrial expansion, such as the export-control problem faced by American firms [2], simply do not exist in Japan. And government-supported projects specifically intended to foster industrial growth are a commonly used government method of supporting



industry, the best-known example being the Super-Speed Computer System being supported by the Government of Japan [28]. To cite the large American R&D investment in defense as an equivalent policy in the United States is a categorical error: this kind of investment is aimed at another objective, and any contribution to industrial expansion is minimal and incidental.

A third myth about Japan is that competition by Japanese firms is the same as competition by the firms of other nations. The difference is what Peter Drucker calls "adversarial" trade as contrasted to "competitive" trade. In competitive trade, a country typically exports the same type of products that it manufactures with the aim of getting a share of a market, as case in point being the reciprocal trade in automobiles between Germany and the U.S. In adversarial trade, the objective is not just a share of the market but the market itself--a case in point being the continuing destruction by Japan of the American and European semiconductor industries. Japanese firms were competing well but not overwhelmingly with the firms of other nations through the mid-1970s. In 1976, MITI (the Ministry of International Trade and Industry) sponsored a VLSI Project that included Japan's six largest semiconductor manufacturers, with the goal of creating the technology for a one megabit chip and thereby giving Japanese firms a dominant share of the market. At that time, the standard was the 16K chip, but through the structural advantage gained from this government-industrial cooperative project Japan now leads the world in this field, with the impending demise of many companies in Europe and the USA. These other nations are now attempting to use cooperative research to regain lost ground, but they fail to understand the scope of the problem, which includes not only an adversarial national industrial policy but also adversarial corporate practices, including "dumping" and "anticipatory pricing," as noted in Section 3.4.

A final example of structural problems is to be found in Japan's "free ride" in defense and foreign aid. Japan's well known limit of 1 percent of its GNP for defense contrasts with some 6 to 7 percent for the U.S. and typically 3 to 5 percent in Europe [11]. Added to this is the fact that Japan does not carry its full share of the foreign-aid burden: whereas the U.S. spends \$800 per capita on defense and foreign aid, Japan spends only \$135. It has been suggested that Japan could begin to carry its foreign aid burden through an Asian version of the Marshall Plan [21], but this would be contradictory to the single Japanese goal of industrial expansion, and thus Japan has not picked up on this idea. These "savings" for Japan effectively lower the tax burden of Japanese firms which can then invest these funds in new product developments.

## 5.2 Government relations.

Specific implications of a partnership economy in which the government has as its primary objective the expansion of industry can be seen in (1) the differing effects of export controls in the two nations, (2) performance of government-supported supercomputer research R&D in Japan but not in the USA, and (3) the closure of the government market in Japan to US firms.

First, Cray Research reports that export controls in the U.S. are costing about 145 days for approval on the average, whereas in Japan the same approval takes less than 30 days. This fact is being used by Japanese firms in marketing their supercomputers: customers are being told by the Japanese that these delays are an unavoidable part of the American offerings but not of the Japanese offerings, and that there inordinate delays are possible, such as the 300-day delay suffered by the University of Stuttgart in Germany in obtaining approval for its Cray computer.

Second, in Japan, vendors of supercomputers are doing research for Government-sponsored supercomputer projects that are set up to benefit the industrial firms, but this is not true in the USA. This has the benefit of providing direct and specific government subsidies for supercomputer research in Japan, whereas the American companies must bear this burden out of their own resources. In other words, the American firms are competing against not just the giant Japanese firms themselves but against the combination of the Government of Japan *and* industrial firms. Some specific cases in point are the Super-Speed Computer Project, supported by \$100 million from MITI, and the Next-Generation Industries Project to develop advanced components needed by Japanese supercomputer firms, among other objectives.

Finally, the de facto [= true in the real world, whether it is true on paper or not] closure of Japan's government market (including their universities) to American supercomputer firms is in contrast to the openness of the American government market for the Japanese. Of the few (seven) American supercomputers installed in Japan, *none* is installed at a government site. Several Japanese universities have expressed interest in acquisition of a Cray supercomputer, but during the procurement process they have been advised that this would cause "political" problems for the funding of their universities if they should actually acquire an American supercomputer, and they have uniformly retreated from such acquisitions [10]. [NOTE: The Japanese will attempt to counter this argument by pointing to the installation of the IBM 3090 high-end mainframe as evidence of the openness of their supercomputer market, but this computer is not considered part of the supercomputer market by either IBM or American supercomputer firms.]

### 5.3 Cultural driving forces.

We have previously pointed out [see reference 8] the effects of the Japanese cultural driving forces in the market place, including their intense work and education ethic, their management style, and their three sacred treasures (lifetime employment, *nenko* reward system [= age priorities], and enterprise unionism).

## 6. SUMMARY AND CONCLUSIONS

### 6.1 Current market status.

- By its large market share, its timing of new product introductions, its strong marketing infrastructure, and its technical leadership, Cray Research has been able to hold off the threat to the supercomputer market by the Japanese competitors.
- As ETA Systems brings its ETA<sup>10</sup> systems to market, their market share will be better protected against incursion by Japanese vendors. However, this is a narrow window in time that could be closed by either new products from Cray Research or the Japanese firms, and the delays being experienced by ETA in bringing this system to market are a serious threat to their very survival.
- Japanese vendors have largely but not completely captured the Japanese market through the traditional "buy Japanese" national bias of the Japanese culture and pressures from the Government of Japan. The Japanese government market is closed to American supercomputer firms. The American vendors been able to place their products only in private Japanese firms.

6.2 Key issues for the future. The main threats to the American supercomputer vendors include the following.

- The vertical integration of the Japanese supercomputer vendors that gives them control over the development of high-performance components and therefore a timing advantage in introducing new generations of supercomputers; the American firms must solve the problem of access to high-performance components in a timely manner.
- The large Japanese firms attempt to "buy" this market away from the small American firms by "anticipatory pricing"; it is difficult to see how to prevent this other than through action by the American government.
- The export control problems being faced by American firms but not by the Japanese firms; this problem must be solved by the American government.
- The closure of the government market in Japan to U.S. supercomputer firms; the American government should assure that American supercomputer vendors find a market in Japan that is as open as is the American supercomputer market.

## 7. REFERENCES AND BIBLIOGRAPHY

1. Karel G. van Wolferen, "The Japan Problem," *Foreign Affairs*, Winter 1986/87, pp. 288-303.
2. Colin Norman, "Academy Panel Blasts U.S. Export Controls," *Science*, Vol. 235, 23 January 1987, pp. 424-425.
3. U.S. Congress, Office of Technology Assessment, *Supercomputers: Government Plans & Policies--A Background Paper*, OTA-BP-CIT-31 (Washington, DC: U.S. Government Print Office, March 1986).
4. Stuart Auerbach, "U.S. Ponders Computer Trade Charge," *Washington Post*, February 5, 1986, pp. F1, F12.
5. Houston Area Research Center (HARC) News Release, January 15, 1986.
6. John W. Wilson, et al., "Is It Too Late To Save The U.S. Semiconductor Industry?" *Business Week*, August 18, 1986, pp. 62-67.
7. P. D. Lax, "Report of the Panel on Large-Scale Computing in Science and Engineering," National Science Foundation report NSF 82-13 (1983).
8. W. J. Worlton, "Japanese Supercomputer Initiatives," *Frontiers of Supercomputing*, N. Metropolis, D. H. Sharp, W. J. Worlton, K. R. Ames (ed's), (University of California Press: Berkeley, 1986).
9. B. Buzbee, R. H. Ewald, and W. J. Worlton, "Japanese Supercomputer Technology," *Science*, Vol. 218, 17 December 1982, pp. 1189-1193.
10. Private communication, Suzanne P. Tichenor (Cray Research, Inc.) to Michael B. Smith (Deputy U.S. Trade Representative), January 9, 1987.
11. S. J. Solarz, "A Search for Balance," *Foreign Policy*, 49 (Winter 1982-83), pp. 75-92.
12. Lewis J. Lord, et al., "The Brain Battle," *U.S. News and World Report*, January 19, 1987, pp. 58-64.
13. Kenneth Fleet, "Treasury joins attack on Japanese barriers," *The Times*, February 12, 1987, p. 21.
14. Kimio Ohno, "Japanese Supercomputers and Molecular Orbital Calculations," *Supercomputer Simulations in Chemistry*, M. Dupuis, Ed., (Springer-Verlag: Berlin, 1986), pp. 49-54.
15. "Japanese Electronics: NEC," *The Economist*, April 12, 1986, pp. 73-76.
16. Willie Shatz and Tom Murtha, "Supercomputers: Planting the Flag," *Datamation*, May 1, 1986, pp. 24-28.
17. David E. Sanger, "Japan's Fast-Computer Plans Falter," *International Herald Tribune*, November 29-30, 1986, pp. 1,13.

18. James A. Martin, "Chip merger concerns U.S.," *Computerworld*, November 10, 1986, p. 126.
19. "Special Report: Japan's technology agenda," *High Technology*, August 1985, pp. 21-67.
20. Peter Drucker, "Doing Business in Japan Isn't 'Just Done'," *The Wall Street Journal*, July 18, 1985.
21. Sam Nakagama, "Let Japan Put Its Money Where Its Might Was," *The Wall Street Journal*, October 18, 1985, p. 30.
22. John Adam, "High technology called key to raising U.S. productivity," *The IEEE Institute*, November 1985, p. 6.
23. W. David Gardner, "Japanese Prepare New Wave of Assaults on U.S. Market," *Information Week*, September 2, 1985, pp. 30-37.
24. Mark L. Goldstein, "Yamamoto's Rising Sun," *Industry Week*, April 15, 1985, pp. 63-64.
25. Sidney Fernbach, "Supercomputer research and development in the U.S. and Japan: an update," *The IEEE Institute*, April 1985, p. 5.
26. Sachio Kamiya, et al., "Practical Vectorization Techniques for the 'FACOM VP'," *Information Processing 83*, R. E. A. Mason (ed.), (Elsevier Science Publishers (North Holland), 1983, pp. 389-394.
27. Charles L. Cohen and George Leopold, "NEC's Latest CPU Keeps Honeywell Line Going," *Electronics*, February 17, 1986, pp. 15-16.
28. Hiroshi Kashiwagi, "The Japanese Super-Speed Computer Project," *Future Generation Computer Systems*, Vol. 1, No. 3, February 1985, pp. 153-160.
29. Hisashi Horikoshi and Rasuhiro Inagami, "Japanese Supercomputers: Overview and Perspective," *Proc. IEEE Int. Conf. on Computer Design: VLSI in Computers (ICCD '84)*, pp. 227-231.
30. Tim Carrington and Robert Greenberger, "Fight Over India's Bid For Computer Shows Disarray of U.S. Policy," *The Wall Street Journal*, February 24, 1987, pp. 1, 24.
31. James A. Martin, "Japan ignoring chip trade accord," *Computerworld*, February 23, 1987, pp. 83, 87.

# Science has added a third tool - computation

---



- Theoretical and experimental science are now joined by computational science
- Not only has computer speed increased by a million fold
- But our understanding of the methods needed to “model” the real world has allowed another million-fold increase in problem solving speed and ability
- As a result computer generated solutions give us insight into areas where experiments are too expensive, impossible, or only useful to verify the computational results

Robert Borcherds UCL



## Historical perspective

---



- **Computing has been a critical element in the Laboratory's programs (especially Nuclear Design) since the start of the Laboratory.**
- **Hardware capacity increased roughly 10 times every 5 years from before 1955 to the early 1970's but the rate has slowed in recent years as fundamental limits are reached.**
- **While manufacturers varied, the largest scientific machines have always been available to LLNL scientists.**
- **Early systems used batch processing, but LLNL introduced interactive use (time sharing), with consequent improvement in design productivity about 1968.**
- **Operation of a new center to serve a national community (MFECC) began in 1973.**

## Trends in computing at LLNL (1985)

---



- Vigorous program in code development in Defense Systems, Physics, Engineering, Lasers, MFE and elsewhere has continued to tax the hardware and systems to their limit.
- We now run almost exclusively on commercial hardware, which has enhanced reliability and the speed with which we can get new systems into production.
- We must still develop our own operating systems but industry trends and increasing use of scientific computing could change that also.
- Efficient use of multiple processors in computer networks is our current major challenge in systems and applications software.




# What is the trend in hardware and software?

---



- **Hardware continues to improve in speed but;**
  - rate has slowed
  - design cycles are long
  - there are fundamental limits
  
- **Parallelism is clearly a possibility for increased throughput but;**
  - it introduces a need for creative improvements at all levels of the problem solving process
  
- **Lack of access to large machines has hindered University research in large scale computing except in limited areas such as magnetic fusion**
  
- **Limited use in the industrial sector makes supercomputers a low volume, specialty product**



# Characterization of each decade



	<u>Components</u>	<u>Architecture</u>	<u>Software</u>
1950	Vacuum tubes Rotating storage	Index registers Parallel I/O	Assembly language
1960	Transistors Core memory	Independent functional units (SISD)	Optimizing compilers
1970	IC K-bit chip	Vector functional units (SIMD)	Automatic vectorization
1980	LSI M-bit chip	Multiple processors (MIMD)	Automatic partitioning
1990	VLSI, JJ Ga-As, HEMT G-bit chip	Massively multiple processors (data flow)	Applicative languages



## Current standards for Class VI, VII, VIII



	<u>Class VI</u>	<u>Class VII</u>		<u>Class VIII</u>	
	Fast uni-processor with vectors 1-6 million words	Fast 4 × multiprocessor with vectors 32-256 million words		Fast 16 × multiprocessor with everything 512-2048 million words	
	<u>Uni</u>	<u>Uni</u>	<u>Multi</u>	<u>Uni</u>	<u>Multi</u>
NET MFLOPS	10	25	100	50	500
PEAK MFLOPS	50	100	300	200	1600


**MFLOPS**

**KNSU**



**Performance  
Indicators** →

APRIL 1986



# Why Multiprocessing?

---



## Performance

- Limitations on single processor performance due to having reached physical limits of
  - Device switching speed
  - Speed of electromagnetic transmission

## Economics

- Find ways to use many low cost VLSI processors and memories

## Assertion

---



**“Within 5–10 years an amount of money, like that needed to buy a present day supercomputer, will be able to purchase essentially arbitrary amounts of computing (measured in raw MFLOPS).”**

- **Whether such machines or systems are supercomputers depends on**
  - **System issues, e.g., I/O bandwidth, memory capacity**
  - **Suitability of the architecture to the problem or *vice versa***
  - **Quality and availability of the software**
  - **How clever we have gotten in parallelizing applications**



# Success at Putting a Problem on a Certain Machine Depends on Suitability at Many Stages of the Problem Solving Process

---



- The problem itself
- Choice of mathematical model
- Choice of discretization method
- Choice of numerical methods
- Arrangement/implementation of algorithms

# Basic Techniques vs. Applications



Applications Area										
	Sparse Linear Systems	Non-linear Algebraic System	Sparse Eigenvalues	Rapid Elliptic Solver	FFT	Multigrid Solver	Stiff ODE Solver	Monte Carlo	Integral Transform	
	1	2	3	4	5	6	7	8	9	10
Lattice Gauge	X			X					X	
Structural Dynamics	X		X	X						
Weather Simulation					X	X				
Circuit Simulation	X		X					X		
Quantum Chemistry	X			X				X		X
Computation Fluid Dynamics	X		X		X	X	X			

Some of these techniques are well known to be suited to new architectures, others may be.

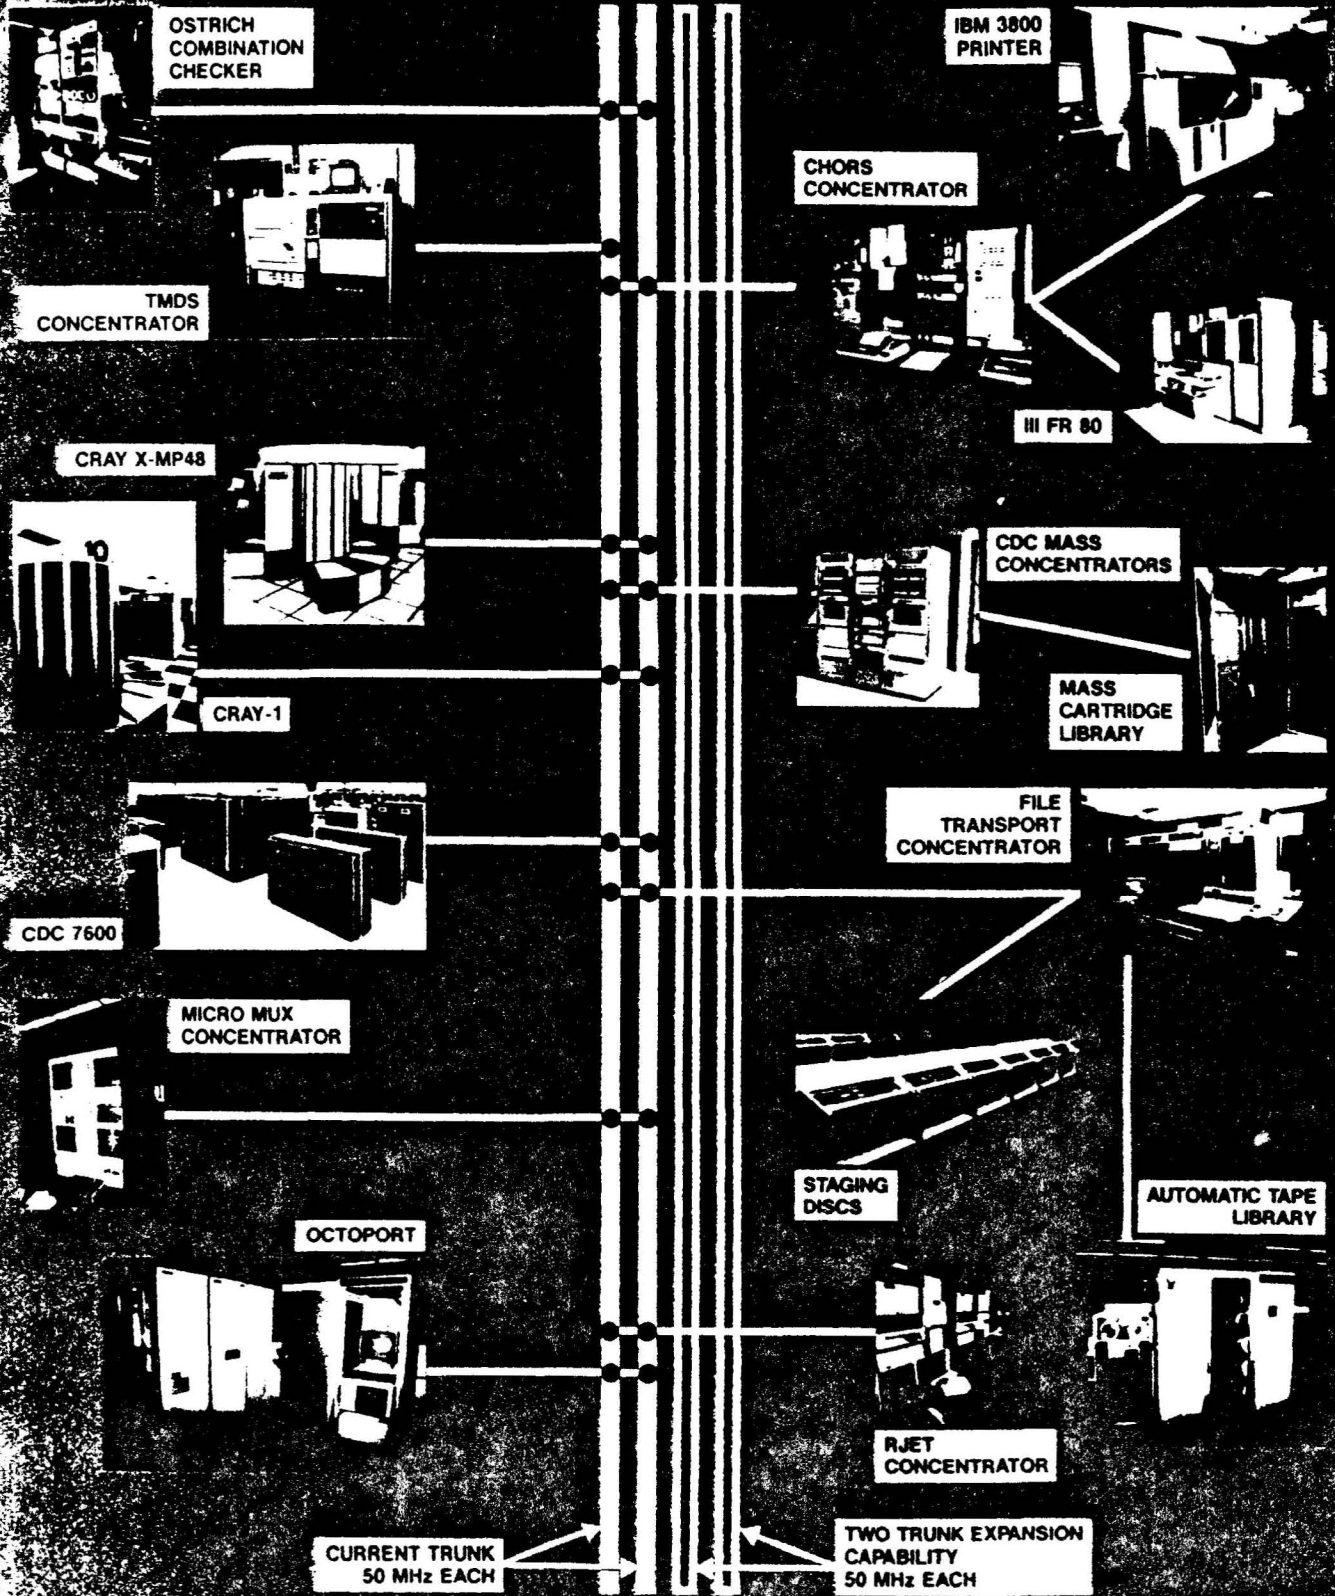
# Performance of Systems Depends on Much More Than CPU Power, Especially for Multiprocessors and Large Applications

---



- Efficient management of asynchronous tasks and efficient interprocessor communication
- Large, efficient memories
- High-speed channels and peripherals
- High-density secondary storage
- High-speed networking technology; for example, protocols and interfaces
- Very high-speed graphics
- Productive environment for software development and management
- Hardware assists for debugging and performance measurement
- Fault tolerance and graceful degradation of multiprocessor systems

# THE OCTOPUS NETWORK



## **Distributed computing**

---



- **Goal: make better use of expensive resources and improve human productivity**
- **Processes can freely communicate, no matter where they are located in the network**
- **Resources are accessed the same, whether they are local or remote**
- **Services can be provided on the best available hardware and moved as required without user impact**
- **New services are easily added - by users or system developers**

# Specialization of Function

---



## Supercomputers

- High speed processing
- Local file cache integrated with central storage
- Network connection
- Local clock

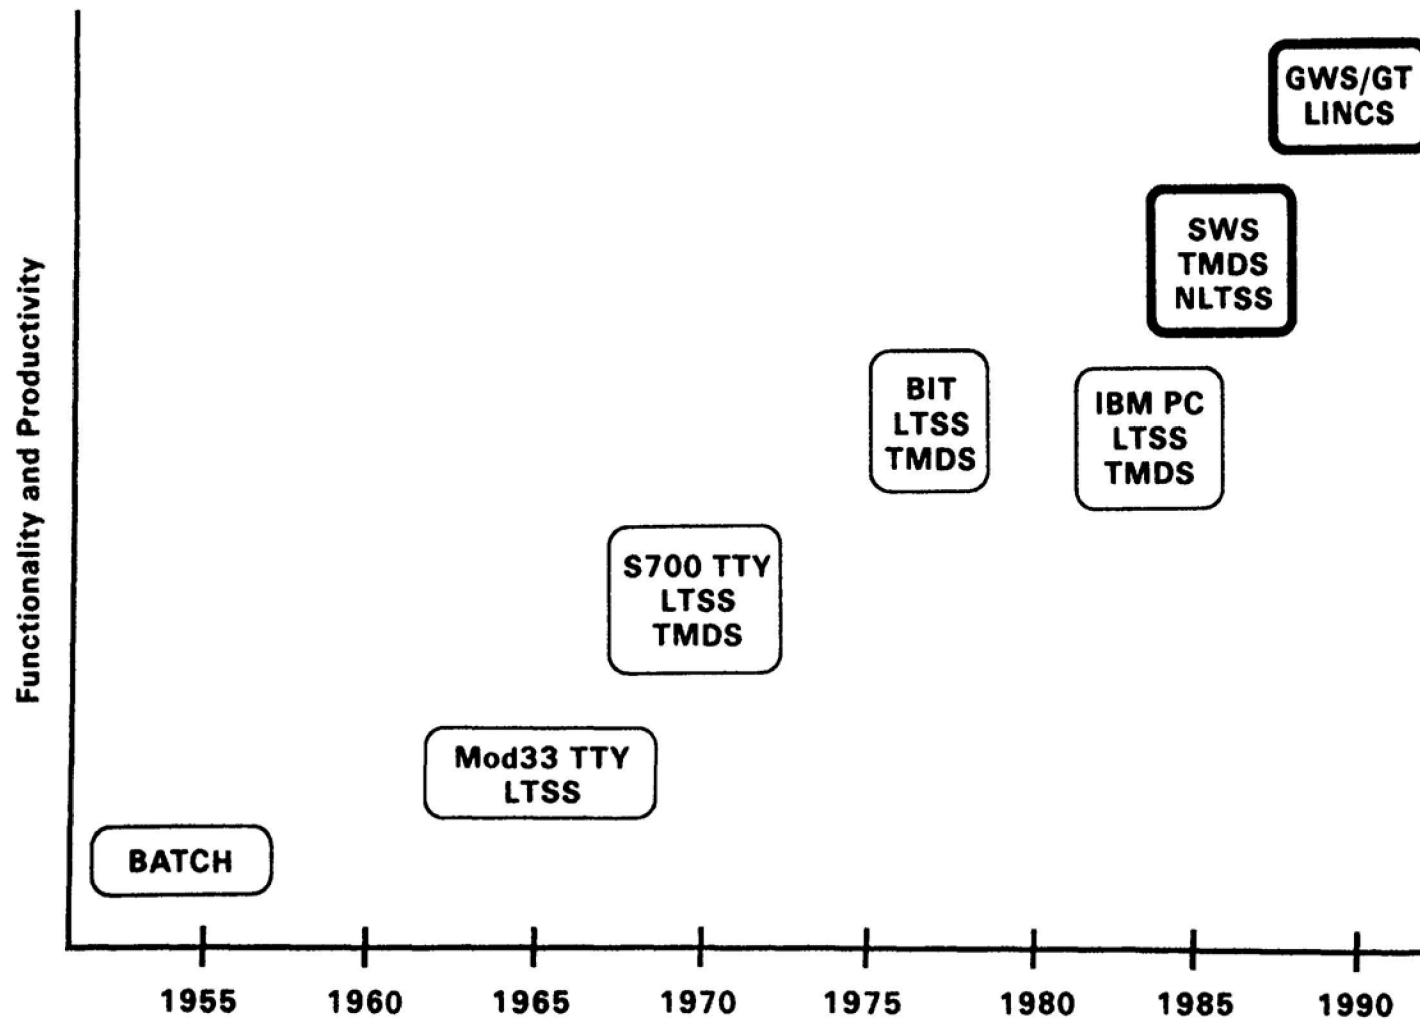
## Workstations

- Command interpretation
- Window manager
- Program development utilities
- Local interactive and graphic applications support

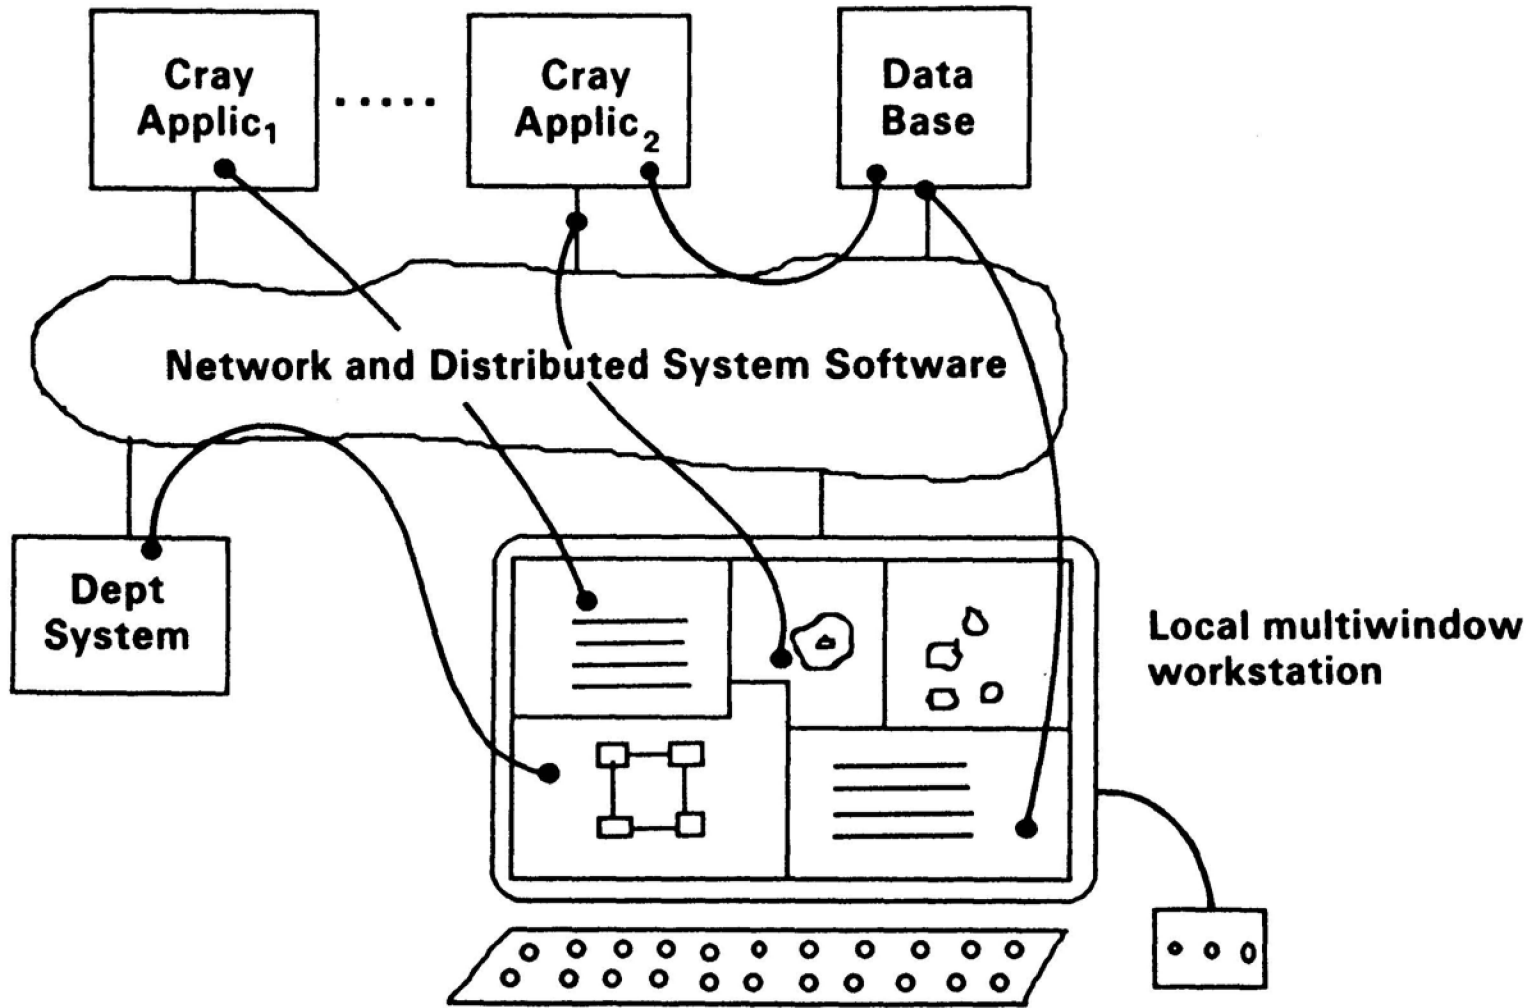
## Shared resources in LCC and departments

- |                                   |                                |
|-----------------------------------|--------------------------------|
| ● Processing                      | ● Time                         |
| ● Authentication                  | ● Mail                         |
| ● Storage                         | ● Data Bases                   |
| ● Printing                        | ● Compilation                  |
| ● Resources Naming<br>(directory) | ● Other software systems, etc. |

# Evolution of the ND User Environment



DR86-01







---

**A Supercomputer is a device that converts  
compute-bound problems into I/O problems**

**Ken Batcher  
Goodyear Aerospace**

## **The computational bottleneck of the late 80's**

---



**The environment of Supercomputer Systems is changing dramatically. To date we have always characterized our computational requirements in terms of CPU performance and memory size.**

**This has led to a large disparity between mainframe performance characteristics and storage (secondary and long-term) performance characteristics.**

**Commercial solutions to correct this imbalance are at least 5–10 years away.**

**We must change our operating philosophy.**

**Worlton and Associates**  
3089 Villa, Los Alamos, NM 87544  
(505) 662-4011 or 662-2724

December 26, 1986

Gordon Bell, Assistant Director  
National Science Foundation  
1800 G Steet, NW, Room 306  
Washington, D.C. 20550

Dear Gordon:

I read with considerable interest your paper, "Preparing for Changing Scientific Computing Environments," and I thought it might be useful if I were to comment on it in some detail.

### 1. THE SCIENTIFIC COMPUTING ENVIRONMENT

I believe we agree that it is most effective for scientific and engineering computing resources to be organized into a hierarchy of technologies and several types of computers. My own perspective is shown in Attachment #1. Along the left side are the four generic information technologies [processing, human interface, storage, and communications], and along the top are three types of resources, defined by the level of sharing: personal resources that are not shared at all [workstations, PCs, and terminals]; mid-range resources that are typically shared by tens of users [super-minis, mainframes, and mini-supers]; and large-scale resources that are often shared by hundreds of users [ideally these are supercomputers]. A major weakness of many scientific computing environments is that they are using mid-range computers as their large-scale resource, where they should be using supercomputers. Using mid-range computers for large-scale problems is a categorical error in the same sense as using personal resources for mid-range problems.

### 2. THE DATA

Some of the conclusions in this document are based on a set of data that has only limited applicability, and it worries me that you are drawing general conclusions from it. There are several problems with the data.

- The performance data comes solely from the LINPACK environment, which has limited applicability as a model of a general-purpose scientific computing environment. Further, you use only the LINPACK data for relatively small problems (100x100 linear systems) rather than also including larger problems (300x300 linear systems, the third table in Dongarra's report), which differ by as much as an order of magnitude.
- The cost data are not normalized to a common configuration. In particular, there are wide differences in the amount of storage and the number of channels. Making cost comparisons of computers whose memories differ by orders of magnitude will surely lead to dubious conclusions. For example, the smaller configurations may *appear* to have equal or even better performance-to-cost ratios than some larger systems, but the smaller systems cannot even run the same class of problems that the larger systems can.

Comments  
12/28/86

Linpac  
≡ Aug.  
H. Wong

~~Worlton~~  
to Dept.

- Thus, the performance and cost data used in this report apply only to those environments with *small problems*, not to the general-purpose class of problems found in the scientific computing environment. General conclusions about scientific computing should not be based solely on small-problem environments.
- With both performance and cost data of limited applicability, the metric of "flp/\$" is doubly suspect.
- Your Table of Computer Characteristics lists "M.s" as being secondary memory size in Megabytes, but you list only 9.6 for the Cray and ETA systems. This puzzles me because 9.6 Megabytes is less than the size of the main memory of these machines. The SSD of the X-MP series has up to 128 Mwords or 1 Gbyte, and even larger versions are in the announcement phase. I suspect there is a decimal point off here.
- Your table includes some introduction dates of '87, '88, '89; if you were going to include proposed machines in this time frame, why not the Cray Y-MP and the Cray-3?

I read the  
to get the...

## 2. THE DIAGRAM

The diagram you include has several problems. First, of course, is the limited applicability of the data on which it is based, as noted above. Second is the rectangular shape of the figures that are used for coverage of the different zones. This imprecision of coverage distorts the comparison of the different types of computers because the upper left portions of these figures are mostly empty. In fact, any family of computers follows a performance-to-cost orientation that is more precisely shown as an ellipse that has its major axis along a 45-degree line--although the exact orientation is manufacturer, product-line, and time dependent. Third, the diagram does not include the results for 300x300 systems, even though the data is readily available in the same report as the 100x100 data. Fourth, "Proposed Single-User Supers" are included that appear to have much better performance-to-cost ratios than any other kind of computers; however, this appearance is flawed by the fact that no comparable proposed systems of other types are shown, including the well-known proposed supercomputers.

I have made a revised copy of your diagram (see Attachment #2), with a number of changes: (1) I have used ellipses, rather than boxes; (2) I have given super-minis and mainframes their own ellipses; (3) I have added an ellipse for "Supers" for 300x300 linear systems from Table 3 of Dongarra's data; (4) I have included an ellipse for "Proposed Supers" running 300x300 linear systems; and (5) I have added a "300x300" ellipse for mini-supers.

I realize that this is still imprecise in some cases, but I believe it is more precise than the original. [I do not have adequate cost data to do a complete revision; if you have a table of cost data for the systems of interest, I would appreciate getting a copy.]

## 3. THE CONCLUSIONS

If we take the new diagram and the caveats noted above into consideration we come to some quite different conclusions than those in the original paper.

- Even if we use them for small problems using algorithms that do not exploit their potential, supercomputers outperform any other kind of computers, and they have a comparable performance-to-cost ratio.

- If you run large problems with appropriate algorithms on supercomputers, the performance-to-cost ratio improves by about an order of magnitude to a level that is not matched by any other kind of system. In other words, the cost per calculation for large problems is lower on supercomputers than on any other kind of computers.
- The mini-super performance improves moderately when you increase the problem size from 100x100 to 300x300 linear systems, but not as much as the supercomputer performance.

#### 4. PAPER MACHINES

There are several computers in the developmental pipeline that should be treated as "paper machines," including the ETA<sup>10</sup> and the Cray Y-MP. The best definition of a paper machine is, I believe: *Any system that has yet to be delivered in significant quantities to customers.* In this context I have developed two useful rules (Worlton's Rules of Paper Machines) to keep in mind:

Rule #1: A paper machine can always outperform a real machine.

Rule #2: Don't believe Rule #1.

It is a logical error to compare something that exists to something that does not exist, as though they were in the same category. Thus, it is an error to compare the ETA<sup>10</sup> to the Cray X-MP; rather, it should be compared to the Cray Y-MP. To understand why this is true, consider the analysis in Attachment 3; I call this analysis "Existence Issues," and all machines still under development face all of these issues. (1) Will it *ever* be completed? Several machines have been brought to an "almost complete" status and have not survived their marketing birth trauma. For example, the Burroughs 8500 was shipped to U.S. Steel, found to be unsatisfactory, and then withdrawn from the market. (2) If it is completed, *when* will it get to market? Delivery dates during the last phase of development are notoriously uncertain. (3) If it is completed, what will the *cost* of development turn out to be? If this is too large, it may sink the company or cause the product to be withdrawn. (4) Assuming that the above issues go well, what will the *performance* be? New supercomputers are also notorious for their hidden bottlenecks. (5) Given all of the above working out well, will the machine be *reliable* in the long term? And finally, (6) if everything else listed above is OK, will the machine still be *relevant* when it is delivered? Relevance has to do with hitting the window of opportunity before it is closed by other products. What other products could close this window of opportunity for the ETA<sup>10</sup>? An announcement by a Japanese vendor of a much faster multiprocessor (the NEC SX-2 already has a 6 ns clock, and their motto is: *Attack!*); an announcement by IBM of 8 PEs and/or a reduction of the cycle time for the 3090; the Cray Y-MP, with a faster cycle time, larger memories, and rich software; etc.

The point is this: Existing machines should be compared to existing machines and proposed machines to proposed machines. We should not blur the critical differences between these categories.

#### 5. BACKGROUND

In this section, there are a few comments that deserve mention from a different perspective. On p. 3, paragraph 4, you say, "Thus, a new class of mini-supercomputer was formed, all of which have better performance/price than the Cray (almost a factor of 2 in the case of the new SCS-40)." I am not sure what you mean by "the Cray." If you mean the Cray-1, you should use a price of not more than \$3.5 million, which is about the price of the used Cray-1; compared to that, the SCS-40 delivered to the San Diego Supercomputer Center has a price of \$0.9 million, which gives it the *same* performance/price as a Cray-1. Further, the X-MP/11 has a price-to-performance ratio of \$4M/1.5 for a cost of \$2.67M per Cray-1 equivalent. To match that a mini-super that operates at 1/4 of a Cray-1 has to be priced at no more than \$2.67M/4 = \$668K, which



puts the X-MP in the same performance/price domain as the mini-supers. Further, often such comparisons do not take the larger memories of supercomputers into account. If you do that, then the economy of scale of the Cray X-MP/416 offers even better memory+processor economics than the X-MP/11. The point is that the common belief that mini-supers offer higher performance-to-price ratios than the supercomputers does not hold up under analysis, especially if we include memory (as we should). A careful comparison of supercomputers and mini-supers must include a qualification about the class of problems for which they will be used, such as the following:

- The absolute cost of mini-supers is lower than that of supercomputers, and for *small-scale problems* their performance/price is comparable to supercomputers.
- The absolute performance of mini-supers is lower than that of supercomputers, and for *large-scale problems*, their performance/price is lower than that of supercomputers.

In the last paragraph on p. 3, you say, "The factor of 5 difference in the speed of the ETA-10 versus a Cray X-MP should open up new problem solution domains." This estimate is off by about a factor of 2. The X-MP has a cycle time of 8.5 ns, which puts the 10.5-ns ETA<sup>10</sup> at a disadvantage of a factor of roughly  $8.5/10.5 = 0.81$  per processor; thus, a 4-processor ETA<sup>10</sup> may well be *slower* than a 4-processor X-MP. Even with 8 processors (which may push ETA delivery times into the Y-MP domain) the ETA<sup>10</sup> may have a factor of approximately  $0.81 \times 8/4 = 1.62$  advantage over a 4-processor X-MP, not a factor of 5. Even the 7 ns ETA<sup>10</sup> with 8 processors will have a hypothetical advantage over the X-MP/4 of only  $(8.5/7) \times (8/4) = 2.4$ , not 5. This remains to be seen, of course, because we are comparing known performance on the X-MP with projected performance on the ETA. Thus, the marketing problem for ETA is this: the 10.5-ns ETA system (which does not yet exist, has only developmental software, and a small market share) will have to compete initially with the 8.5-ns X-MP (which exists, has rich software, and a large market share). The 7-ns ETA system will have to compete with the Cray Y-MP (with a presumably faster clock) whose delivery will occur in the same time frame.

On p. 4 you say, "For example, a slower machine is likely to be used more interactively and results of the computation viewed constantly to avoid unnecessary work." Why do you consider interactivity to be relevant only to slower machines? The supercomputers in use at Los Alamos, Livermore, the Sandia National Labs, and many other sites including the NSF sites at San Diego and University of Illinois are all using supercomputers in an interactive mode, with the "results of computation viewed constantly to avoid unnecessary work."

## 6. THE "STRETCH" CONCEPT

I gather that by "stretch" you mean the amount of run time that the user of a lower-cost, slower machine would gain relative to time on a supercomputer due to the cost differential of the machines. The problem with using this concept to describe scientific computing environments is that it is based on a categorical error: you are comparing low-priority work on a supercomputer with high-priority work on a smaller computer, and this is an incomplete description of scientific computing environments. Attachment #4 attempts to make this clear. Here we show the space defined by a range of performance and a range of priorities, with four zones of interest labeled A to D.

- Environment #1 is employed by those sites using the "bankpoint" concept developed at Livermore that allows users to manage their own priorities. This method is used in the CTSS system employed by the NSF Centers at San Diego and Illinois, as well as various DOE sites: the user can "stretch" his allocated time by running at low priority or accelerate his response time by running at high priority. The advantages of this environment include the dynamic nature of user-controlled priorities and the common operating environment for all of the user's work. The user need not shift to a different operating environment for his low-priority work, he just changes his priority.
- Environment #2 is employed by those sites using distributed computing, in which an organization may have a low- or mid-range machine such as a VAX on site under their own control that is used as a front-end to a supercomputer. They can run small-scale tasks on the VAX and then submit large-scale runs to the supercomputer. Los Alamos implements this environment through its XNET system, with some 100 VAX computers attached to its computing network. Cray Research implements this environment through "station software" in which any of about a dozen types of systems can serve as front-ends for the Cray computers.
- Environment #3 describes the stretch concept, where a site chooses between getting low-performance computers and using them in high-priority mode, or sharing a supercomputer in low-priority mode.

The point is that the stretch concept is an incomplete and inadequate descriptor of the choices available for developing a scientific computing environment. It fails to recognize the critical roles of allocation and priorities in the management of supercomputers. The ideal scientific computing environment has all three of these options available for users; some of the NSF sites (San Diego and Illinois) do that and some of the DOE sites do, as well.

## 7. THE ROLE OF THE NSF SUPERCOMPUTER CENTERS

I am concerned by your apparent change in policy on the role of the NSF supercomputer centers. They were originally formed to provide large-scale computing resources for scientists and engineers working at the frontiers of science, technology, and engineering. Universities often can and do manage the acquisition and availability of workstations and mid-range computers on their own; however, in many (and perhaps most) instances they cannot manage the availability of supercomputers on their own, and that is the justification for the NSF supercomputing initiative.

Now, however, you seem to be diverting these centers away from their original role of supporting research in science and engineering toward a role that is essentially computer science: serving as beta test sites for "all new systems", supporting "experimental machines," and trying out "all forms of computation." To add this new set of responsibilities to the centers you will either have to provide them with new funding for a staff of computer scientists or you will inevitably weaken their ability to support research. If you are able to offer new money, that is good; if you are planning to dilute their original role with new responsibilities without new funding, then this is certainly a major change in policy, one that needs to be discussed in the community and a consensus reached before the policy is put into place.

*entire*

*performance benchmarks in report*

*Supercomputer Industry goal:*

*sig gain will come through OS understanding*

Thank you for the opportunity to comment on this paper. Building a consensus before taking action is, I believe, one of the fundamental principles of good management, for it is consensus that gives us the support we need for our actions. Only dictators do not need to manage by consensus. Failure to manage by consensus is the primary cause of management disasters, including the current Iranian-Contra crisis. Building a consensus takes time, but not nearly as much time as cleaning up the messes we create when we try to take actions without a consensus.

Very sincerely yours,



Jack Worlton

wjw

cc: E. Block, Director, NSF  
J. Connally, Program Director, OASC  
R. Ewald, CRI  
Derek Robb, CRI  
File: NSF



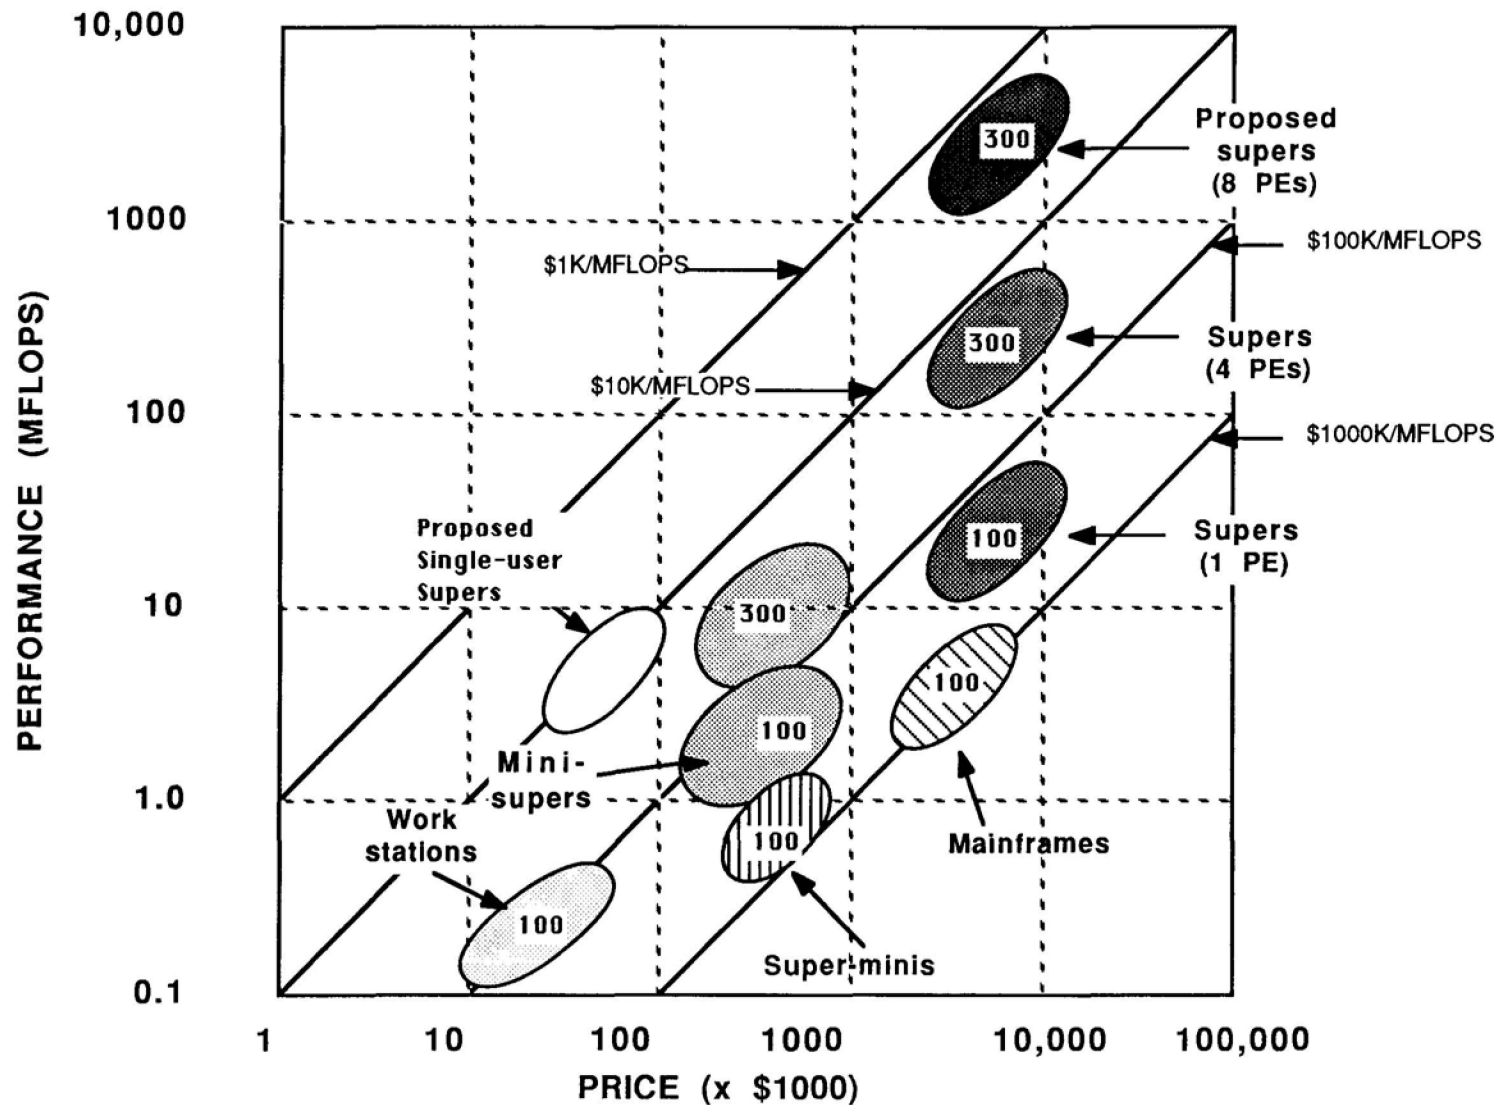
# THE SCIENTIFIC COMPUTING ENVIRONMENT

## RESOURCES

TECHNOLOGIES	LARGE-SCALE	MID-RANGE	PERSONAL
PROCESSING	Supercomputers	Mini-supers Mainframes Super-minis	Workstations PCs and terminals
HUMAN INTERFACE (I/O)	High-speed graphics and printers	Medium-speed plotters and printers	PC graphics and printers
STORAGE	Common file system	Local disk systems	Floppies and hard disks
COMMUNICATIONS	Site networks, LANs, WANs	Site networks, LANs, WANs	Site networks, LANs, WANs

NOTE: WAN = Wide-Area Network; LAN = Local-Area Network

# Comparative Linpack Performance



Data source: J. Dongarra, ANL Tech Memo No. 23, 8/5/86, and G. Bell, NSF.

SC-p/c by wjw  
10/29/86: wjw

# EVALUATING SUPERCOMPUTERS: Existence Issues

- $P_{\text{existence}} = \prod P_i$  , where

$P_1 = \text{Completion}$

$P_2 = \text{Schedule}$

$P_3 = \text{Cost}$

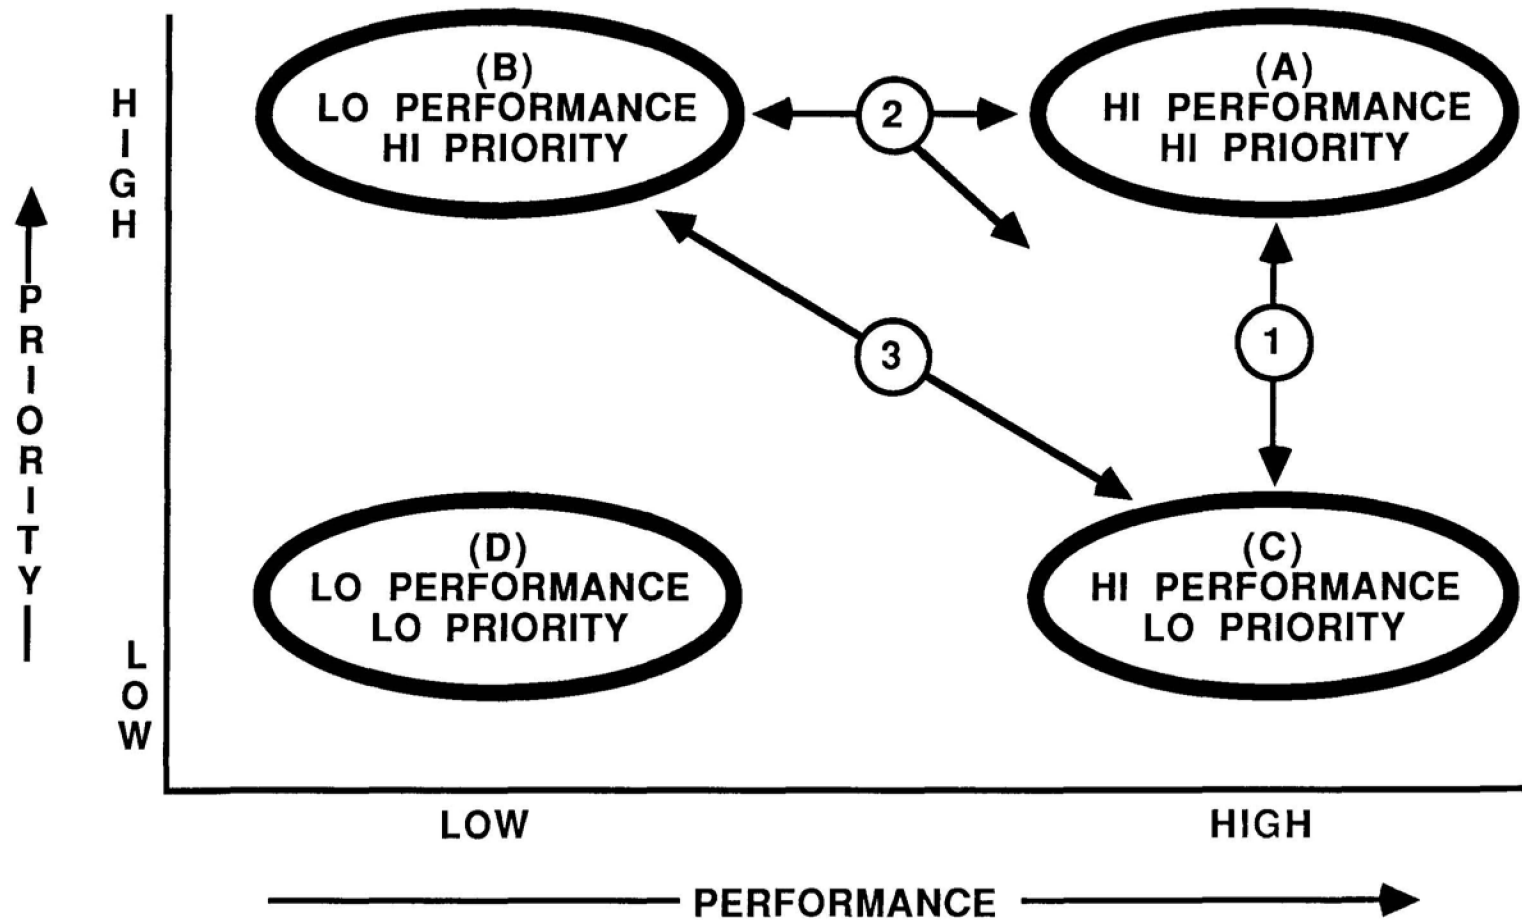
$P_4 = \text{Performance}$

$P_5 = \text{Reliability}$

$P_6 = \text{Relevance}$

- $P_{\text{existence}} \leq \min (P_i)$
- $P_{\text{existence}} = 0.26$  if  $P_i = 0.8$  for all  $i$ .
- $P_{\text{existence}} = 0.53$  if  $P_i = 0.9$  for all  $i$ .

# SCIENTIFIC COMPUTING ENVIRONMENTS: PERFORMANCE VS PRIORITY



# ISSUES AFFECTING EXPORT CONTROLS FOR SUPERCOMPUTERS

by  
Jack Worlton  
Laboratory Fellow  
Los Alamos National Laboratory

The purpose of this report is to analyze certain issues that have arisen in formulating policies governing the controls necessary for the export of supercomputers. This is a *draft for comment*, and if there are additional issues that need to be discussed or comments on the issues discussed herein, please send them to the author at 3089 Villa, Los Alamos, NM 87544.

ISSUE #1: *Can usage of a supercomputer reveal information about its design other than that already available in the open literature?*

This issue is based on a categorical error--the failure to distinguish between *learning* and *using*. In some activities, we learn by doing, such as riding a bicycle, but in others we must learn a great deal before we can undertake the activity. This latter category includes science, technology, and engineering in general, and supercomputing in particular. We must learn *how* to use a supercomputer before we can use it, and this is done at either the Fortran level or the Assembly Language level. High-level languages such as Fortran 77 are designed to be machine independent, so we learn very little if anything about the architecture of the machine from these languages. On the other hand, Assembly Languages are designed to allow access to the detailed characteristics of the machine contained in hardware description manuals, and it is at this level that the details of the machine are revealed. However, the specifications of the Assembly Language are in the open literature, so anyone wishing to learn about characteristics of a machine through Assembly Language or machine specifications already has this information available to them.

Some details of the machine design are not revealed through usage, of course. By using a supercomputer--or any other computer--one does not gain access to essential information on its technical design, including components, timing, power, and cooling. Thus, we conclude the following:

*Usage of a supercomputer merely confirms information that is already contained in the published specifications of the machine and of the Assembly Language. Technical characteristics of the supercomputer such as the components used, critical timing characteristics, power, and cooling cannot be learned through usage of the machine.*

*Computer*

ISSUE #2: *What benefits are obtained from running a problem on a supercomputer that are different from running it on another type of computer?*

The answer to this issue is contained in the relationship among three variables: response time, complexity, and execution rate, as follows:

$$\text{Response time [Time/problem]} = \frac{\text{Complexity [Operations/problem]}}{\text{Execution rate [Operations/time]}}$$

The complexity of a problem is determined by how many operations must be executed to complete it; current large-scale problems have complexities of  $10^{11}$  to  $10^{12}$  total operations. If we now divide this measure by the execution rate, we arrive at the time to solve the problem. For example, if the complexity of a problem is  $10^{12}$  and we execute that problem on a computer that has an execution rate of  $10^8$  operations per second, then response time would be  $10^4$  seconds, or 2.8 hours; on the other hand, a problem with a complexity of only  $10^{11}$  operations run on the same computer would have a response time of  $10^3$  seconds, or about 15 minutes. The sources of complexity are determined by the following four-factor formula:  $C = G \cdot T \cdot V \cdot A$ , where

$$C[\text{ops/problem}] = G[\text{points/problem}] \cdot T[\text{time-steps/point}] \cdot V[\text{variables computed/time step}] \cdot A[\text{operations/variable}].$$

Thus, complexity grows as we use finer grids of mesh points and more time steps in order to obtain more detailed information; this is necessary as the devices we are studying become more complex. Complexity also grows as we add more detailed physics to the problem to make it correspond more closely to the real world; we have to compute more variables, and the algorithms specify more operations per variable.

Attachment #1 shows the relationship among complexity, execution rate, and response time as a "nomograph." In this figure, any straight line across the three vertical scales shows a correct relationship among the variables. Thus, line A shows that  $10^9$  operations executed on a computer running at  $10^6$  operations per second would yield a response time of  $10^3$  seconds [about 15 minutes], whereas line B shows that  $10^{12}$  operations executed on the same computer would yield a response time of  $10^6$  seconds [about 278 hours].

We can now use this nomograph to study the issue at hand. Suppose American computers are some 10 times faster than Soviet computers; how does that affect the class of problems that can be solved on them? The nomograph shows two things: (1) the same problems can be run on the two types of machines, with the difference being that the response time on the slower computer is 10 times longer than on the faster computer; (2) the same response time can be achieved on the two types of machines if the complexities of the problems being run differ by a factor of 10. Thus, any problem that can be solved on an American supercomputer can be solved on a Soviet computer if the Soviets are willing to wait 10 times as long for the answer; this is an advantage to the American researcher, but it isn't as though the Soviet researcher couldn't get the answer at all, for he certainly can.

Would a Soviet researcher be likely to run a critical (secret) calculation on a supercomputer in another nation? Suppose the Soviet researcher could use either one hour per day of time on a supercomputer in another nation (that to him is not secure), or 10 hours per day on his own machine that *is* secure and readily available. Which would he choose? His own machine, of course: to do otherwise could easily compromise the secrecy of his project and gain almost nothing.

The above discussion assumes that the computational problems being solved on the two types of computers have the same complexity, but this might not be true. There are many algorithms that solve the same problem but have widely differing numbers of operations. Thus, it is possible for the Soviet researchers to apply complexity-reduction methods to the application

being encoded before the computer program is written, thereby reducing the number of operations that must be executed. Keep in mind that the Soviet Union has world-class mathematicians. In other words, the American computational position relative to the Soviets is not determined solely by hardware characteristics alone; rather, it is determined by a combination of hardware, system and application software, and mathematics. It isn't necessarily true that Americans have an advantage in all of these areas. Thus, we come to the following answer to Issue #2:

an inside  
in a possible  
in a  
sanitary  
env.

*Run time on a supercomputer is quantitatively different, but not qualitatively different, than run time on any other computers. Thus, with so little to gain, and so much to lose, it is unlikely in the extreme that critical Soviet calculations would ever be run in another national environment.*

**ISSUE #3:** *How much time would a Soviet researcher have to use on a supercomputer in another nation to obtain a significant advantage over using his own machines, and can this amount of usage be readily detected?*

The answer to this issue can also be quantified. There are about 8000 hours per year available on a computer. If we assume some application at hand is critical to the Soviets, they would devote some significant percent of a computer to it, say 25 percent or 2000 hours per year. Assuming the factor of 10 advantage to the American computers, this corresponds to 200 hours on an American supercomputer, so anything less than 200 hours of unauthorized supercomputer time would be useless to the Soviets, because they can get that equivalent computing time on their own computer. Thus, we are faced with the question of whether we can detect the unauthorized usage of several hundred hours of computing time. Accounting for the time usage on a computer is a "zero-sum game," i.e., all time must be accounted for. The unauthorized user of a large amount of time is faced with two choices: (1) use the time (somehow) without any account being charged, which would make the total out of balance and therefore be detected; or (2) use other users' charge accounts, in which case the users would detect that they are being overcharged. Small amounts may slip through the accounting system--seconds, minutes, or perhaps even an hour now and then--but certainly not tens or hundreds of hours.

But the problem for really complex systems is even larger than this. A rule of thumb for the design of a nuclear weapon is that it requires about a full year of supercomputer time--about 8000 hours, so if the system under design is of this scale of complexity, the problem for the unauthorized user is not just how to get away with several hundred hours but several thousands of hours of unauthorized time without the accounting system detecting it.


Thus, the answer to this issue is:

*To gain any significant advantage by using a supercomputer in another nation, a Soviet researcher would have to obtain hundreds to thousands of hours of time per year on a foreign supercomputer, but this level of usage would easily be detected by the normal accounting checks and balances.*

  
12/24/86

# NOMOGRAPH of Computer Simulation

COMPUTER RESPONSE TIME (SEC/PROBLEM)	SUSTAINED EXECUTION RATE (OP/SEC)	PROBLEM COMPLEXITY (OP/PROBLEM)
--	---	---------------------------------------





## Datamation Interview

- Paper on changes
- Talk to /usr/group

There's a hierarchy... I would like it compatible. POSIX and standard Networks are very important. This is becoming an important issue, especially as we begin to do Visualization

### Some concerns

- range and diverting money to support smaller machines. No danger, but one might ask -- How do you buy the mid-range machines?
  - > Mini-supers appear to be as cost-effective, given that you're only a factor of 4 down in power. Centers require 35-75 today. A mini super with 1/16 of the power can usually be operated with a few people.
  - > My commitment. "I believe the NSF funded centers should exist with the largest possible supercomputers, forever."
- stopped the Phase I centers as planned. This is especially important in light of the fact that major universities are making commitments to scientific computing: Texas, Berkeley, Ohio State, Michigan... as well as Minnesota who has been a leader with a whole cadre of CDC's and Cray's
- Push to go to Unix. POSIX is a standard. I have in mind a computing environment patterned after the homogeneous environment of VAX/VMS, but that is not instruction-set compatible, based on UNIX. I don't believe we should be developing an O/S (CTSS). Also, we have to "visualize"
- Have gotten involved with other agencies, especially DOE and NASA about committing to UNIX and insisting that the machine work before we pay for it... and that CDC provide help with a second machine to ease the crunch.
- We need more sharing and support from states and industry (eg. Cray)
- getting computer science involved is necessary... today we have none
- Do we have enough users? Are we training enough users? Are we working on big enough problems... have we escaped the VAX mentality.

### The program:

- 4000 users; 200 universities
- supercomputing **power** has gone from essentially 0 before the program to 5/2, to 102 (assuming 80 by IBM's)/38 (4 xmps, 1 cray 1, 2-205's, and a 3090/400>600) today. We see another factor of 4 within 3 yrs.
- Universities are all buying their own too. Surprise to me that places like MIT and Stanford haven't so far.
- Lots of exciting results from helping design the America's cup winner to computational chemistry, to being able to simulate electronic structures.

Stuart Rice, Chemist, U. of Chicago, NSB Member  
IBM ACIS Forum on the Physical Sciences

"I have in mind a networked system ... graphics workstations and local supporting intermediate computer and ultimately connects to a supercomputer, with provisions of special devices...

1. Distribution of computer resources distorted by the use of "funny money" ... cash and credit ... Workstations come from grants, supers are "free" ... intermediate machines are indispensable and current funding patterns will have to change if ...
2. Dramatic advances in hardware haven't been matched by advances in algorithms and operating systems. ... parallelism is "chicken and egg"
3. The scientific community has become rather inflexible with respect to use of operating systems. ... don't use particular machine features...
4. ... the scientific community has not been as imaginative as it might in thinking about the uses of computation in research."

## **Advantages of Common Environment\***

- 1. better support and consulting overall**
- 2. user experience cumulative so users reach higher level of knowledge**
- 3. shorter start-up on new machines**
- 4. easier to plan for/invest in long-term projects (large codes)**
- 5. larger amount of software available (3rd party, users, etc.)**
- 6. encourages vendors to diversify hardware**
- 7. protects medium and small vendors (investment in OS, etc.)**
- 8. Easy interchange with CS community (if Unix) *Attracts/encourages/facilitates use of modern vector, multi-processor computers***
- 9. *Encourages competition, avoids lock-in of supers, ... workstations through portability***
- 10. *Multiple center use requires no re-learning and permits load-sharing***
- 11. *Multiple-class use is necessary, desirable and simplified (supers... workstations)***

\*Cornell Director's meeting, *GBell comments*

## Disadvantages of Common Environment

1. How to get to there from here

*(staff & user retraining)*

2. Lack of vendor support

*(from CDC, DEC, and IBM but not Cray)*

3. Stifles creativity & innovation unless evolutionary path provided

*(innovation in operating systems, sans human interface a'la MAC, has virtually stopped; standards evolution process must support rapid evolution)*

4. Trade-off performance for generality

*(no evidence)*

5. *Generic and batch supercomputer functions*

*(e.g. checkpoint) must be added to Unix*

6. *Risk of domination of computer industry by Japanese manufacturers*

## Requirements for a Common Environment (in order of achievability)

1. emulation of a "standard interface"
2. office functions (mail, text, etc.)
3. networking (telnet, ftp, etc.)
4. applications packages (Linpack)
5. program development tools (awk, sed)
6. compilers (*and language dialects*)
7. program execution environments and  
system resource management;  
*parallel processing primitives*
8. graphics
9. libraries
10. databases

## **Level of Functionality**

**(in order of achievability and importance)**

- 1. Workstations with common environment and bridges to specific computing engines**
- 2. front-ends as time-shared equivalent**
- 3. supercomputers**
- 4. specialized machines (databases ...)**

## **The Real Sticky Wickets**

- 1. Migration path – must convince the users, the staff, and the vendors**
- 2. How to keep innovation and creativity alive and allow diversity**
- 3. How to allow performance optimizations within a standard framework**
- 5. How to choose/specify the common environment**

## Advanced Computers in Univ. Res. 1988 (est.)

Computer Pwr	Number	Power.Cray equiv.
Supers		
Cray 48/2 8	8	64
CDC/ETA 20	2	40. (ETA est.)
IBM 3090 2	50	100
Minisupers .67		
	50	33
Superminis & .1		
Mainframes	500	50
Workstations .02		
	4000	80



LASL Conference on supercomputers sometime before August 20, 1989  
Start by commenting on earlier assumptions:

1. I guess I disagree with the whole tenor of the conference: scientific computing can only be done on supercomputers. Computing is a completely substitutable commodity like transportation. People will go where time is the most cost-effective. Even with free Cray time users didn't bite. The VAX 780 proved it. (see slide about where technical computing is done)
2. Disagree with Hecker: Significant increase in use. The market looks flat as the audience narrowly defines supers.  
# Crays = # users = constant. Some growth through replacement of minis that were used in technical computing because DEC got out of the market.
3. Hardly a rapid increase in demand for supers. Read Business Week: Cray revenues are flat, rumored to be putting out a machine to compete with Convex. Convex is growing and is at the 250 level, TMC is about 40M. A billion or so. Hardly large, it feels like we're spending that much in R&D. Give it to the Japs.
4. Supers are a symbol of competitiveness like Harley Davidson.
5. Supers are the key to competitiveness. I don't believe this at all. Using them may be. It surprises me that Cray uses its own computers to design machines. Would feel better if they had workstations.
6. Demand is insatiable. This should be true if you charge nothing. Even at NSF, where the price was right the demand wasn't there.
7. K was right. Big question is the economic one. Companies died because they had poor computers for the most part... or ran out of fuel \$ or didn't attend to fundamentals of running a business or of marketing. No amount of policy or gov't spending is going to expand the supers market or legislate competence.
8. Nelson had a good question: What would you do if you were starting from scratch to design a complete environment and customer had 20 m to spend. (If I were IBM I wouldn't spend it on traditional, non-scaleable computers.)
  - some mC, until we get these scaleable mP's straightened out (I would never build one of these, I only build mPs that can be used)
  - a small mP server using the i860 successor (a Gflop in a VCR box)
  - a hot 860/486 workstation family for visualization and most computation.
8. Japanese are 10' tall. I agree. They have the gold, will, are willing to do the dirty work of manufacturing machines, technology base, have resources, we have the ideas for free.
9. Dick Clayton said it all. We're here because we like to have fun. I enjoy building machines for this community because the appetite is insatiable.
10. Unclear what will help the fundamental problem: stimulating understanding the opportunity so that users will get more committed. My current bottom line is:  
Nothing! Wait until the PC gets supercomputing built in so that every user has to worry about parallelism.

**Supers are limited by:**

1. The switch. S/C design is a spreadsheet exercise. Pick the price, allocate the switch and memory. Through in the processors.
2. Nets. Centralized computers require communication nets. Nets are really bad.
3. Human interface, interactivity and visualization.
4. Poor price performance and slow evolution of technology. (Even Stardent put about 600 computers out in a year. This amounts to 60-120 YMP processors or 8-16 large Y's. Though in a free processor to color pixels with each one.
5. Sociology of central vs distributed things that are easy to buy.
6. Commitment to a perhaps unachievable clock. Unclear how this gets built? V/Ulsi is a much better approach. Not trying to make water run uphill.