

# The Future of High Performance Computers in Science and Engineering

*A vast array of new, highly parallel machines are opening up new opportunities for new applications and new ways of computing.*

## Gordon Bell

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 the most diverse and exciting one yet. Some of the research accomplishments have been stimulated by DARPA's Strategic Computing Initiative (SCI). Other progress is a result of the evolution of understanding the technology of multiple vector-processing computers (i.e., supercomputers). However, without the scientific base, creative talent, understanding and demanding users, and infrastructure to design complex VLSI chips, these innovative machines would not be possible. A variety of established and newly formed companies have organized to exploit the new technology. Table I lists the number of companies building high performance computers for control and artificial intelligence applications, and traditional supercomputers for scientific and engineering applications.

The impressive gains in performance and the 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.

Three kinds of computers, distinguished by the way they are used, are emerging:

- General purpose computers are multiprogrammed and can handle a variety of applications at the same 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 com-

puters can achieve an increase in performance or performance/price of a factor of 10 or so.

- Applications-specific computers solve only one problem such as speech or image processing. By binding the applications in hardware and software at design and construction time, an increase by a factor of 100 to 10,000 in performance or performance/price is possible.

Figure 1 shows the alternative forms of the computer structures being considered in the preceding section. Figure 2 shows various approaches to building high performance computers and the estimates of performance versus time.

## MAINLINE SUPERCOMPUTERS FOR SCIENCE AND ENGINEERING

The main line of scientific and engineering supercomputer development follows the Cray formula: the fastest clock; a simple pipelined scalar instruction set; a very complex, pipelined vector processing instruction set; and multiprocessors. Software is evolving to automatically vectorize and parallelize so that the increase in the number of processors is often transparent to users.

The supercomputer has been defined three ways by supercomputer architects and engineers [11]:

1. The fastest machine with the most memory available.
2. Ability to solve numerically intense problems. Such a computer is typically characterized using the Livermore Loops or Linpack benchmark execution rates measured in floating point operations per second.
3. The fastest computer in a particular, named price class with both the highest scalar and vector processing performance.

The three classes of supercomputers are also characterized by price: The Cray or "true" supercomputer at

TABLE I. Companies Building High Performance Computers and Supercomputers.

Kind of computer	On market	Developing	Dead	Recent <sup>1</sup>
Supercomputers	2	2?	3	5
Mainframes (with vector proc.)	5	?	1	2
Mini-supers	5	2?	4	10
Graphic Supers	2	?	0	2
Total Supers	14	4	8	19
Array Processors	8	?	4	6
Massive Data Parallel	2	1?	?	3
Multiprocessors	6	2?	4	11
Multicomputers	16	?	2	18
Total (including supers)	46	7?	18	57
Superminis	7	1?	2	3
RISC-based computers	7	?	1	5

<sup>1</sup> Started after 1983.

\$1 million to \$20 million; the mini-supercomputer at \$100,000 to \$1 million; and the graphics supercomputer at \$100,000.

In 1989 personal supers costing between \$15,000 and \$50,000 will be introduced. Such machines are based on Intel's 80860 microprocessor, for example, which operates at 40 Mhz and delivers roughly 10 megaflops on the Linpack benchmark.

In this article, a supercomputer is defined as a computer that:

- Is used for numerically intense computations.
- Has the highest scalar and vector processing capacity with the appropriate primary and secondary memory resources.
- Falls into one of the supercomputer price classes

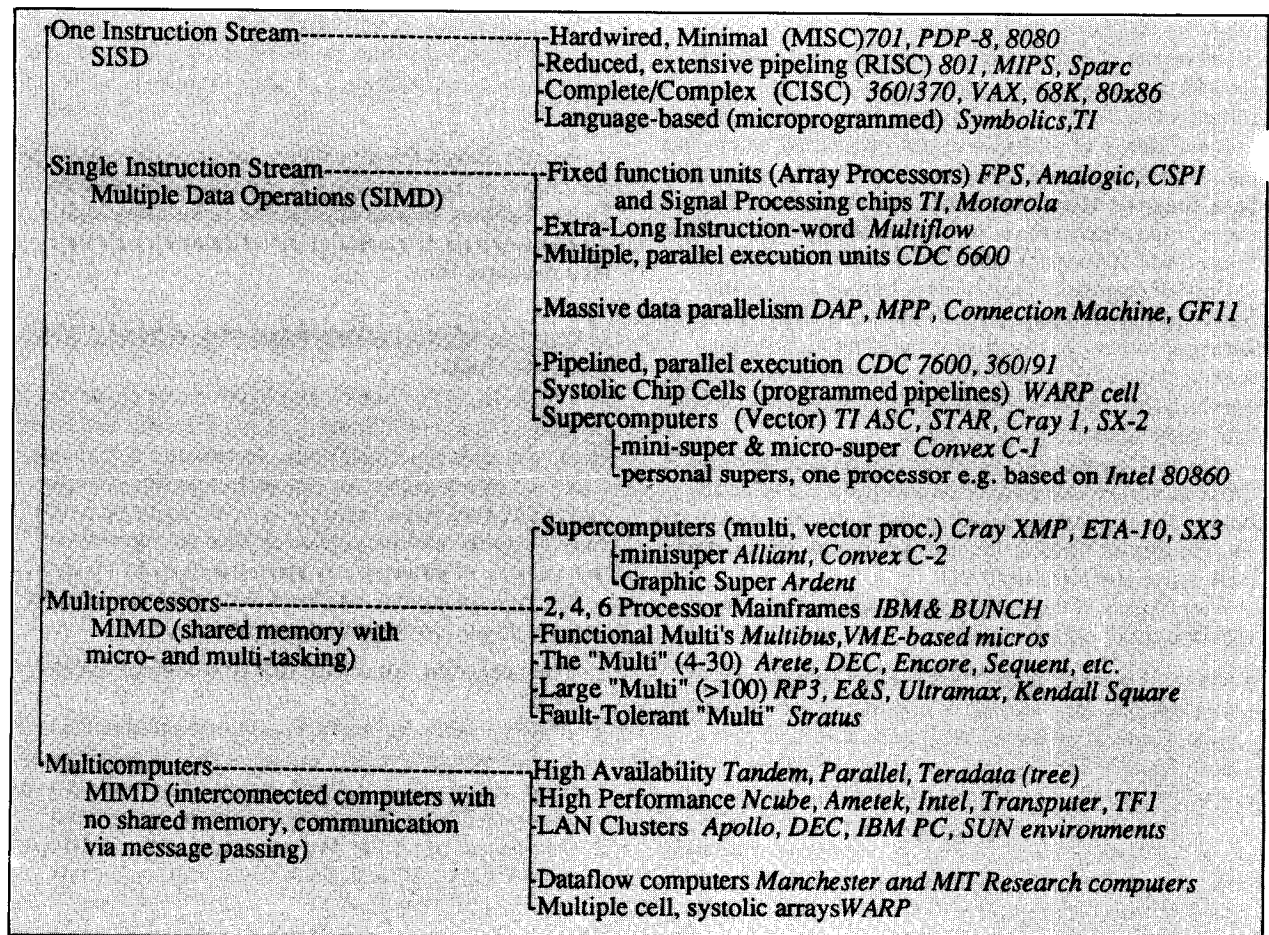


FIGURE 1. Simplified Taxonomy of Various Computer Structures for High Performance Computing and Supercomputing




FIGURE 2. Current and Projected Operation Rates in either Instructions Per Second or Floating Point Operations Per Second for Microprocessors, Multiprocessors (mP), Supercomputers, Multicomputers (mC), and the Connection Machine (CM)

(i.e., super, mini-super, graphics super, personal super).

- Can supply the computational resources (i.e., processing, primary and secondary memory, graphics, and networking) to a single user or set of users in one day that could be obtained using the largest shared supercomputer.

The Japanese computer industry's first supercomputers aimed at having the fastest uniprocessors as demonstrated by NEC's SX-2, Fujitsu's VP series, and Hitachi's S-820-80. These uniprocessors provide very fast single stream execution, but the Cray XMP still has more aggregate throughput. Through its parallelizing compiler, the Cray YMP is the world's fastest supercomputer [6], and is likely to remain so until the Japanese start building multiprocessors. In April 1989 NEC introduced the SX-3, which it believes will be the world's fastest computer. Scheduled for delivery in 1990, it is rated at 22 gigaflops using four processors, each operating at 5.5 gigaflops, and a 2.9 ns clock.

In 1990, Cray Research plans to introduce the Cray 3 (16 processors at twice the speed of the Cray 2), to be followed in 1992 by the Cray 4. The Cray 4 operates at a clock rate of 1 ns and delivers 128 gigaflops, using 64 processors. Figure 3 shows the evolution of clock speed, number of processors, and aggregate computing power for Cray computers over a 30-year period beginning with the CDC 6600.

Getting peak power on a single program will require the user to do some reprogramming to exploit the multiple processors. Each year one Gordon Bell Prize acknowledges the fastest program execution on a traditional supercomputer. In 1988 the prize went to the Long Range Global Weather Modelling Program at NCAR, which was parallelized manually and operates at over 400 megaflops (versus a peak of 840 megaflops) on the Cray 416. Given the relatively small memory and the growth in application program size, researchers are beginning to use the XMP in parallel mode to take full advantage of its processors. In 1989, the award went to Vu, Simon, Ashcraft, Grimes, Lewis, and Peyton [5], who used an 8-processor Cray YMP operating at 1.68 gigaflops (versus a peak of 2.7 gigaflops, or 5.5 times the uniprocessor rate) to solve a large statistics problem using a finite element model.

### DISTRIBUTED SUPERCOMPUTING

Given the interest in supercomputing in government, academe, and industry it is worth looking at the range of opportunities for different styles of computing from the highly centralized or regional supercomputer to a fully distributed computing environment using personal supercomputers. Just as distributed computing using LAN-interconnected workstations and personal computers has become an alternative to mainframes, fully distributed personal supercomputing is likely to play a similar role with respect to traditional supercomputers. However, since the central approach is so strongly subsidized by providing "free" supercomputer time to the scientific community, it is unlikely that




FIGURE 3. Past, Present and Future Projection of the Clock Speed, Number of Processors and Performance of Seymour Cray-designed CDC and Cray Research Computers

distributed supercomputing will evolve quite as rapidly.

### Supercomputing Policy

With the entry of the Japanese into the supercomputer market, supercomputing has become an issue of national pride and a symbol of technology leadership. While Japan now builds the fastest uniprocessors, the multiprocessor approach provides more throughput, and with parallelization, more peak power. Table I indicates that a large number of new companies have started up to build high performance computers since 1983. Both Steve Chen, formerly of Cray Research, and Burton Smith, formerly of Dennelcor, have started such companies. A recent report by the Office of Science and Technology Policy urges the adoption of a government initiative for high performance computing, including a National Research Network [7].

When the Cray 1, the first modern supercomputer, appeared, university users and researchers were using a highly distributed approach using the VAX 780. Several thousand VAX 7xx-series computers were running scientific and engineering applications, and before the Cray XMP was introduced, only one Cray 1 was available at a university. The run-time for the VAX could be up as much as a factor of 100 times the Cray for a stretch factor of 100. The acquisition costs for processing, measured in millions of floating point operations per second per dollar, were about the same. VAX allowed comparatively large programs to be run easily using its large virtual memory.

With the emergence of the Cray XMP in 1983, the performance, capacity, and cost-effectiveness of supercomputers was increased by factors of 1.5, 6, and 3, respectively. New VAX computers with adequate power and performance/price were not introduced to compete with the Cray XMP. Only in 1984, after the XMP was already established, did mini-supercomputers with comparable performance/price emerge. By 1984 a strong national initiative based on the NSF Bardon and Curtis report had begun [2]. It resulted in the establishment of the NSF's Advanced Scientific Computing (ASC) and its five National Computing Centers. The centers are now completely institutionalized into the NSF infrastructure at a budget level of approximately \$60 million in FY1989. The centers and the 5 to 10 percent of NSF researchers who use them have made a very strong case for government-funded supercomputing centers.

State and federal governments are very receptive to large centers. Several states are building supercomputer centers to serve their university systems. Similarly, many corporations believe a supercomputer is required for competitiveness. For example, Toyota has more supercomputer power than all of the U.S. automotive industry. At least one U.S. research laboratory believes that a supercomputer is essential for recruiting.

### Performance and Cost-Effectiveness

There appears to be no economy of scale across the supercomputer classes. Lower cost, higher production

**TABLE II. Comparison of Central Cray YMP and Distributed Titan Graphics Supercomputer**

Factor	Cray YMP 832	Titan 24
Price	20	0.12
Processors	8	2
Mwords of memory	32	4
Drystones (integer-oriented)		
KDhrystones/s single processor	25	23
Dhrystones/s/\$ multiprogrammed	.005	0.383
Whetstones (scalar floating point)		
MWetstones/s single processor	35	6.5
Whetstones/s/\$ multiprogrammed	14	108
Linpack (100 × 100)		
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 × 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	N/A	50

Note: The time stretch factor for Titan relative to the Cray YMP is given in parentheses.

quantity components (processor, memory, peripherals, etc.) result in an inherent dis-economy of scale in performance/price for smaller-scale applications. This is demonstrated in a comparison between the Cray YMP and Ardent's Titan Graphic Supercomputer, both of which were introduced in 1988. On the other hand, having a large central facility is critical for certain large programs and databases.

Table II gives the purchase price, performance, and performance/price for several benchmarks run both sequentially and in parallel for the two approaches. Observe that for the purchase price of the YMP, one could have 166 graphics supercomputers in a highly distributed fashion for project, departmental, and even personal use. The comparison ignores operating costs, which in the case of a central computer are quite visible. In the distributed approach operations costs (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. Well-funded, large centers provide a complete and balanced set of facilities including large primary memories of gigawords as well as being able to handle tens or even hundreds of gigabyte files, and archival files for collaborative science. However, few centers operate at this level. The "right" environment for a user depends on need. Ideally, most users would have their own distributed, cost-effective personal supercomputers for smaller applications with extensive, 3-D graphics and access to central supercomputers via fast networks for running large community programs and accessing large databases.

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 run-time on the Cray YMP. Also associated with each benchmark is the cost-effectiveness or performance/price (e.g., megaflops/second/dollar) of the YMP versus the Titan. The range of results is comparable with an analysis of Titan and the Cray XMP by Misak and Lue at the Institute of Supercomputer Research [8] where for scalar and vector loops, the Cray was faster by about a factor of 5 and 10, 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 the Titan, indicating that while the YMP is still faster for utility programs, it is not cost-effective by over an order of magnitude.

At first glance, the small Linpack case seems irrelevant to supercomputing. However, the average speed which the Cray XMPs run at various large computer centers has in the past been equal to about 25 megaflops/second, or the speed of the Linpack  $100 \times 100$  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 eight processors and has once again become the world's fastest computer. 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 of around 10 to 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 1,000 hours per year, while large users get an hour a day, and average users an hour a week.

By using the peak speeds that are obtained only 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. Thus, we see the importance of parallelization to increase speed on a super in order to provide significantly more power than a user could get in a single day on a personal super.

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. Current networks operating at T1 (1.5 megabits per second) are inadequate for applications requiring extensive computation and visualization such as molecular modelling, computational fluid dynamics, finite element modeling, and volume data visualization [4]. Future supercomputers may have embedded rendering hardware to provide both cost-effective and truly interactive graphics supercomputing.

## FUELING HIGH PERFORMANCE COMPUTERS

The future of the conventional uniprocessor appears to be a simple RISC-based architecture enabling high speed execution and low cost, one-chip (i.e., microprocessor) implementation. Such a processor is able to track semiconductor process technology and is inherently faster than a multiple chip approach. Current microprocessor performance is comparable to that of the IBM 3090 mainframe, which has historically evolved at 14 percent per year. By the end of 1989, the performance of the RISC, one-chip microprocessor should surpass and remain ahead of any available minicomputer or mainframe for nearly every significant benchmark and computational workload. By using ECL gate arrays, it is relatively easy to build processors that operate at 200 Mhz (5 ns clock) by 1990. One such company, Key Computer (recently purchased by Amdahl Corporation), is doing just this to build very fast uni- and multi-processors. Prisma is building a GaAs-based computer using the Sun Sparc architecture. By adding an attached vector processor, users can see a very bright picture for scientific and engineering computation in workstations and simple computers.

The projected evolution of the leading edge one-chip POPs (plain old processors) using the RISC approach is given in Table III. Unlike the historical leading edge clock evolution, POP clock speed has evolved a factor of 5 in four years (1.5 per year) because the processor is on a single chip. Shifting to ECL can give an aggregate speed-up of a factor of 20 over a six-year period (1.65 per year).

## USING LOW-COST, FAST RISC MICROS

The very fast CMOS and soon-to-come ECL microprocessor will compete with every other computer and be a good component for both multiprocessors and multi-computers. Also, the micros can be used in a redundant fashion to increase reliability and build what is fundamentally a hardware fault-free computer.

## Multiprocessors

The next generation, high-performance micros are designed for building multiple microprocessor computers, or "multis" [3]. Thousands of multis are now in operation, and will probably become the mainline for traditional time-shared computers and smaller workstations for the next decade. However, given the speed and simplicity of POPs, users may ask Why bother with so much performance? By 1990 workstations with four to ten 20-mips processors attached to a shared bus in a multi configuration that sells for under \$50,000 may

TABLE III. Past, Presented, and Projected Clock and Performance for Leading Edge, One-Chip Microprocessors

Year	Clock (mhz)	PkMips <sup>a</sup>	Mflops <sup>b</sup>	Mflops with vector unit
1986	8	5	1	—
1987	16	10	2	16
1988	25	16	5	25
1989	40	25		
ECL shift				
1990	80	50	10	100
1992	160	100	20	200

<sup>a</sup> Mips for millions of VAX equivalent instructions per second. VAX 11/780 = 1 million.

<sup>b</sup> Mflops for millions of floating point operations per second using Linpack 100 × 100 benchmark [6].

exist. With only a small incremental design effort, all computers can be implemented as multis; 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 has also been 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 use Cray-style multiprocessor designs, 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 for highly vectorized applications since automatic compilation or parallel constructs to use a large number of scalar processors don't appear to offer the speed of the vector approach for these applications. Furthermore, it is difficult to build very large multiprocessors as cheaply. Hence, multicomputers have been demonstrated and have gained utility for user-explicit, parallel processing of single problems at supercomputer speeds.

The most ambitious multiprocessor being introduced in 1990 for both scientific and transaction processing is by Kendall Square Research. Each 64-bit processor operates at 40 megaflops and has an associated primary memory of 32 megabytes and an 80 megabyte/second I/O channel. The combined system with 480 processors operates at 19.2 gigaflops with 16 gigabytes of memory.

Given the ease of building very high performance multiprocessors based on POPs, a major transition is

likely to occur in the traditional mainframe and minicomputer industries. Using this multiprocessor approach, several computers that execute programs at over 100 million instructions per second and cost less than \$500,000 have entered the market. Such a structure provides 2 to 4 times the power of IBM's largest mainframe at 1/20th its cost. The transition is predicated on the fact that the performance/price discontinuity will cause users to consider switching new and even existing programs to the new micro-based environment and away from mainframes and minis on what is used as a "code museum" to run existing programs.

### Hypercube-Connected Multicomputers

In the early 1980s, Seitz and his colleagues at Caltech developed a large, multicomputer structure known as the hypercube. By using commodity microprocessors, each with private memory to form a computer, and interconnecting them in a hypercube, grid, or switching network, each computer can pass messages to all other computers. Today's multicomputers provide substantially higher performance using faster interconnection networks for message passing than their first generation ancestors, making them more widely applicable. For a particular application, a factor of 10 in performance/price over mainline supercomputers has been observed.

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

The European multicomputer counterpart to the hypercube is based on Inmos' Transputer computer node. A transputer is a processor with a small on-chip memory and four, full duplex interconnection ports operating at 20 megabits/second which are used to pass messages to other transputers, especially when they are interconnected in a grid. Several companies are building general purpose, multicomputers by connecting a large number of transputers together. The transputer is proving especially useful in systems to build applica-

TABLE IV. Multicomputer Generations [1]

	First 1983-87	Second 1988-92	Third 1993-97
Nodes			
MIPS	1	10	100
Mflops scalar	0.1	2	40
Mflops vector	10	40	200
Memory (Mbytes)	0.5	4	32
Number of nodes <sup>a</sup>	64	256	1,024
Message time (us) <sup>b</sup>	2K	5	0.5

<sup>a</sup> Typical system. Maximum system is roughly four times larger.

<sup>b</sup> 100 Mbyte packet.



tion-specific systems for everything from communications to robots.

Over a quarter of a billion dollars has been spent to build the 15 present generation multicomputers. Users have undoubtedly invested a comparable amount in programming. It is unclear how successful these machines have been as measured either by establishing healthy, growing, profitable, and hence, sustaining businesses or by providing users with computational power to solve unique problems. Current systems only approach supercomputer power while requiring users to reprogram; thus, users trade-off lower operational costs for a higher initial investment.

Multicomputers can become an established structure only if they provide power that is not otherwise available and the performance/price is at least an order of magnitude cheaper than traditional supercomputers. The next generation of multicomputers (e.g., NCUBE's 8K nodes operating at 2.4 megaflops) with power in the same range as the 20-gigaflop supercomputers available in 1990 will be important in establishing multicomputers. The DARPA-funded Intel multicomputer with 2K nodes that operate at 80 megaflops each is equally important.

## NEW RESEARCH MACHINES

The following machines have emerged from DARPA's Strategic Computing Initiative (SCI) or other research in basic computer science.

### Systolic Processors

Kung's work at Carnegie Mellon University on systolic arrays is beginning to pay off and arrays are being applied to a variety of signal and image processing tasks in military applications. The 10-cell WARP operates at an average of 50 percent peak for the problems of interest (e.g., speech and signal processing). This provides 50 megaflops for \$350,000 (142 flops/second/dollar) and is available from General Electric. 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 1991 time frame. Using the chip, a small board could compute at roughly 100 megaflops. It would not be unreasonable to expect this component to sell for \$10,000 or provide the user with 10,000 flops/second/dollar. While the initial product was developed for a special purpose, the ability to use the WARP for a range of applications is improving with better understanding of the compilation process.

### Text Searching Machines

Several companies have built specialized text and database machines. Hollaar at Utah has built and field tested machines for very high speed, large database text searches. The result to date is that inquiries 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.

## The Connection Machine

The Connection Machine grew out of a research effort at MIT by Danny Hillis, which resulted in the establishment of Thinking Machines Corporation. Several machines are installed in a variety of applications that require parallel computation, including text searching, image processing, circuit and logic simulation, computational fluid dynamics, and finite element method computation. The current CM 2 model is a uniprocessor with 64K processing elements. It has up to 1/2 gigabytes of primary memory, and operates at speeds up to 10 giga-floating point operations per second. Thus the CM 2 is the supercomputer for a number of applications. While the Connection Machine operates on one problem at a time in a mono-programmed fashion, it should be able to be multi-programmed.

## Evolution of the Array Processor

Multiflow Corp. was started up based on Fisher's work at Yale on a compiler to schedule a number of parallel (7 to 28) execution units using an extra wide instruction word. 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 mini-super speeds and costs half as much. One feature of this approach is that a compiler can exploit a substantial amount of parallelism in existing "dusty Fortran decks" automatically. Cydrome, now deceased, built a similar product using ECL technology that provided higher performance and better performance/price using a similar architectural approach.

## EXPERIMENTAL RESEARCH MACHINES 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 that 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 highly unlikely that any computer specialized for a particular language will be able to keep up.

### University of Illinois Cedar Multiprocessor Project

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

**Very Large Multiprocessors**

Three SCI projects—BBN's Monarch, Encore's Ultramax, and IBM's RP3—all explore the size and utility of large multiprocessors and provide over 1,000 mips in sizes of 1,000 at 1, 128 at 16, and 512 at 2 mips performance, respectively. None have vector processing, and hence may not 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**

This machine is a scaleable, highly parallel, SIMD architecture that provides the highest performance for a variety of graphics processing tasks such as solids rendering under varying and complex lighting situations.

**AT&T's Speech and Signal Processor**

A large number of signal processing computer chips arranged in a tree-structured multicomputer configuration provides over 250 gigaflops on 32-bit numbers. The machine fits in a rather small rack, and the resulting number of flops/second/dollar is nearly one million. The machine came, in part, from Columbia University's tree-structured multicomputer work and is part of DARPA's SCI.

**IBM Research GF11 and TF1**

GF11 was designed specifically for Quantum Chromodynamics calculations and is a SIMD computer providing 11 gigaflops. TF1 has a goal of achieving 1.5 teraflops using 32K 50-megaflop computers connected via a large, central switch to pass messages among the computers.

**Special Algorithm Attached Processors**

Several computers for special applications in chemistry, computational fluid dynamics, genome sequencing, and astronomy have been built or are in development. Each provides factors of 100 to 1000 over comparable hardware using a general purpose approach.

**COMPUTER TECHNOLOGY, RESEARCH, AND TRAINING NEEDS****Faster Circuitry**

University research in high speed circuitry, interconnection, and packaging is virtually nonexistent. Very high speed processors require much better interconnection and packaging density.

**Mass Storage**

No radical improvements in size or speed are in progress to keep up with the processing developments that come from VLSI. A project that would couple 1,000

one-gigabyte drives in a parallel fashion to provide an essentially random access to a terabyte of memory and use a variety of specialized architectures is feasible. A project at Berkeley is exploring this approach [10]. 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. The 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, under NSF sponsorship, recommended a significant research and development program in visualization [9]. Today's supercomputers are capable of generating data at video rates. In order for humans to interpret data, it appears that the best way is to use direct couple, 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 mini-supercomputers 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 array 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. An alternative to the Hyperchannel LAN, which operates at a peak of 50 megabits/second, is needed. The next generation must be a public standard. A combined fiber distributed data interface (FDDI) and a nonblocking, public standards-based switch that operates at 100 megabits/second using fiber optics seems like a necessary first step that could evolve over the next several years.

**Wide Area Networks**

Intermediate speed (45 megabit) and fiber optic (multi-gigabit/second) switch and network development and research are not occurring rapidly enough for computer networking. The networking dilemma is well-defined [4, 7]. These networks are badly needed to interconnect the plethora of local area and campus area networks. Today, many 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/second to connect with the 1,000 to 2,000 academic, industrial, and government research organizations. By making a system that could be used for both computers and communications, the two disciplines and industries could begin to become synergistic rather than antagonistic.



### Memory Address Limits

The address limit of 32 bits on most of today's computers begins to be a severe constraint for every configuration but multiple computers. For example, a solids data set could easily have an array of  $1,000 \times 1,000 \times 1,000$  elements, requiring a 33-bit address.

### Dataflow as an Alternative Computational Model

Arvind's group at MIT has progressed to the point where it is building a dataflow computer that might 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 a specialized architecture will be able to keep up with the main line.

### Neural Computing Networks

Various efforts aimed at highly parallel architectures that attempt to simulate the behavior of human processing structures continue to show interesting results.

### Changing the Programming Paradigm

It is necessary to change the programming paradigm to get the highest performance from 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. Table V summarizes the main line of computational models and the corresponding near-term computer structures that support the model.

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 models in Table V may not be the best or even adequate to express parallelism. For now, however, we should build on what we know, learn from experience, and research alternatives.

### THE TERAFLUP COMPUTER BY 1995

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

1. A 4K node multicomputer with 800 gigaflops peak or a 32K node multicomputer with 1.5 teraflops.
2. A Connection Machine with more than one teraflop and several million processing elements.

Both types of machine require reprogramming according to either the message passing or massively parallel data with a single thread of control programming model. However, to exploit the evolving supercomputers with 16 to 64 processors, users will have to do extensive reprogramming that will use a combination of micro- and multi-tasking and message passing.

Today's secondary memories are hardly adequate for

TABLE V. Summary of Main Line Computational Models

Computation model	Supporting computer structures
Vector processing	Supercomputers (one processor)
Message-passing (coarse-medium grain)	Workstation clusters on a LAN Multicomputers (e.g., hypercubes)
Micro-/multi-tasking (fine grain)	Shared-memory, multiprocessors Shared-memory, multiprocessors
Massive data-parallel single thread of control	SIMD (e.g., Connection Machine)
Dataflow	Special dataflow multicomputers, multiprocessors?, and SIMDs?

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 per job quite high, thus diminishing the performance/price advantage. Based on current applications, it is likely that either approach will be more than 10 percent as efficient as a general purpose, multiprogrammed, shared memory computer. The cost of such machines will be comparable to the supercomputer of the 1990s, which is likely to be over \$50 million.

### APPLICATIONS CAN EXPLOIT THE OPPORTUNITY

While it is difficult to predict how the vast increase in processing power will affect science and engineering generally, its impact in the following specific areas is clear. Given the current level of training, however, it is unclear how rapidly applications will evolve to exploit the enormous opportunity available in significantly faster and more cost-effective computers.

#### Mechanical Engineering

Computers are being used in the design of mechanical structures ranging from automobiles to spacecraft, and cover a range of activities from drafting to analyzing designs (e.g., crash simulation for cars). Designers 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 significant improvement in mechanical design. Under this design paradigm, every facet of product design—including the factory to produce the product—is possible without prototyping. Within the next decade the practice of mechanical engineering could be transformed provided companies seize the opportunity. For starters, better product quality could result from the new technology, but the big impact comes from drastically reduced product gestation periods and the ability to have smaller organizations, which also contributes to product elegance and quality. A similar change in chip and digital systems design has occurred in the last decade, whereby almost any chip or system can be designed and built in about two years.

#### Biochemistry, Chemistry, and Materials

In molecular modeling and computational chemistry,

the design of molecules is now being done interactively with large-scale computers.

#### **Large-Scale Scientific Experiments Based on Simulation**

A dramatic increase in computational power could make a range of system simulation involving many bodies (e.g., galaxies, electron interaction at the atomic level) feasible. Nobel laureate Ken Wilson characterizes computer simulation as the third paradigm of science, supplementing theory and experimentation. This paradigm shift will transform every facet of science, engineering, and mathematics.

#### **Animation**

Large-scale computers can compute realistic scenes, providing an alternative to traditional techniques for filmmaking.

#### **Image Processing**

Various disciplines including radiology rely on the interpretation of high resolution photographs and other signal sources. By using 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 applications that will be feasible for the PC. For example, Ardent's graphic supercomputer, Titan, which currently sells for about \$100,000 is an excellent model of what will be available in 2001 for less than \$6,000 (assuming a continued price decline at the rate of 20 percent per year). Every home will have an almost unlimited laboratory to conduct experiments.

#### **ROLE OF THE FEDERAL GOVERNMENT**

Dis-economy of scale continues to exist and favor small machines. This suggests that personal supercomputing, like its counterpart in ordinary computing, will migrate to a distributed approach provided decisions are made on some rational or economic basis. Large regional computers are difficult to access effectively using today's limited networks, especially for interactive visualization. NSF's Advanced Scientific Computing (ASC) program successfully supplies 5,000+ researchers (representing only a few percent of the scientific community) with an average of 1 hour of supercomputing time per week. Smaller supers such as mini-supers or graphics supers could easily supply researchers with the equivalent of 10 to 40 hours per week. Most agencies, however, have no means to support smaller, more cost-effective computers unless the prices are at workstation levels that can be supported by research grants. By not using small, dedicated computers that are under the direct control of the researchers who use them, users are deprived 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.

While highly specialized computers offer the best performance/price and operate at supercomputer speeds, they cost more than workstations. AT&T's speech processor that carries out 1/4 tera floating point operations per second on 32-bit data is an excellent example of the gains possible by applications specific hardware and software. Again, agencies are unlikely to risk building such specialized computers, nor does the community have the right combination of computer scientists, 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 provided the problem is rich in data parallelism and users are willing to reformulate and reprogram their applications. NSF's research and computing directorates support only traditional supercomputing. Only now are agencies with the greatest need (DOD, DOE, HHS, and NASA) beginning to examine the Connection Machine for computation. The need is clear if my own thesis is correct about achieving the highest performance.

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 occur, will work on nearly all high performance computers including the Connection Machine, multicomputers, and multiprocessor supercomputers such as the Crays. The shared memory, micro- and multi-tasking models 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 to bring it nearer the state of the art is an important development that will begin to allow general use and 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 and that autotasking compilers are evolving to take advantage of this in a transparent fashion to a limited degree. The bad news is that not all applications can be converted automatically. Users are not being trained to use such machines in an explicit fashion. No concerted effort is in place covering the spectrum from training to research. This will require a redirection of resources.

#### **A ROLE FOR THE COMPUTER AND COMPUTATIONAL SCIENCE 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 community can continue to ignore applications in computational science. Alternatively, it can learn about the various forms of parallelism supported by the evolution of mainline supercomputers and the onslaught of new, highly parallel machines by being involved with the applications, by using and understanding the new structures, by writing texts, training students, and carrying out research. For example, no current texts about programming are based on the mainline development of the supercomputer. Understanding may enable work on automatic programming systems to analyze and rewrite programs for these computational models. As a minimum, understanding will greatly facilitate exploiting both the evolution and the radically new machines.

#### REFERENCES

1. Althas, W.C. and Seitz, C.L. Multicomputers: Message-passing concurrent computers. *IEEE Comput.* 21, 8 (Aug. 1988), 9-19.
2. Bardon, M., and Curtis, K. A national computing environment for academic research. *NSF Working Group on Computers for Research*. National Science Foundation, July 1983.
3. Bell, C.G. Multi's: A new class of multiprocessor computers. *Science* 228 (Apr. 26, 1985), 462-467.
4. Bell, C.G. A national research network. *IEEE Spectrum* 25, 2 (Feb. 1988), 54-57.
5. Brown, J., Dongarra, J., Karp, A., et al. 1988 Gordon Bell Prize. *IEEE Software* 6, 3 (May 1989) 78-85.
6. Dongarra, J.J. Performance of various computers using standard linear equations software in a Fortran environment. *Argonne National Laboratory Technical Memorandum*, 23 (Feb. 23, 1989).
7. Office of Science and Technology Policy. *A Research and Development Strategy for High Performance Computing*. (Nov. 1987).
8. 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* 2, 1 (Aug. 1988), 8-11.
9. McCormick, B.H., DeFanti, T.A., and Brown, M.D., Eds. Visualization in Scientific Computing. *Comput. Graphics* 21, 6 (Nov. 1987).
10. Patterson, D., Katz, R., and Gibson, G. A case for redundant arrays of inexpensive disks, RAID. In *Proceedings of the ACM SIGMOD International Conference on the Management of Data* (June 1988).
11. Perry, T.S. and Zorpette, G. Supercomputer experts predict expansive growth. *IEEE Spectrum* 26, 2 (Feb. 1989), 26-33.
12. Smith, N.P., Ed. *Supercomputing Review* 1, 1 (1988).
13. Smith, N.P., Ed. *Supercomputer Review* 1, 4 (Apr. 1989).

**CR Categories and Subject Descriptors:** B.0 [Hardware]: General; C.0 [Computer Systems Organization]: General; C.1 [Computer Systems Organization]: Processor Architectures; C.3 [Computer Systems Organization]: Special-purpose and Applications-based Systems; [Computer Systems Organization]: Performance of Systems; C.5.1 [Computer System Implementation]: Large and Medium (mainframe) Computers—*super (very large) computers*; J.2 [Computer Applications]: Physical Sciences and Engineering; K.1 [Computing Milieux]: The Computer Industry; K.3.2 [Computers and Education]: Computer and Information Science Education; K.4.1 [Computers and Society]: Public Policy Issues

**General Terms:** Design, Human Factors

**Additional Key Words and Phrases:** Circuit technology evolution, parallel processing, supercomputers and other high performance parallel computers

---

#### ABOUT THE AUTHOR:

**G. GORDON BELL** is Vice President of Research and Development at the Ardent Computer Company. His research interests include the contributions to and the understanding of the design of interesting, useful, and cost-effective computer systems, and the organizations that build them. Author's Present Address: Ardent Computer Company, 800 West Maude, Sunnyvale, CA 94086.

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the ACM copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Association for Computing Machinery. To copy otherwise, or to republish, requires a fee and/or specific permission.

