

Shopping for Supercomputers

In the market for supercomputing? Look beyond Cray supercomputers, IBM mainframes, Alliant minisupers and Digital minis. Instead, discover the power of superworkstations.

BY C. GORDON BELL AND GLEN MIRANKER

The supercomputer, despite its high political and high technology profile, is giving way to a new set of distributed technologies built around workstations as the chief means of technical computing in the 1990s.

The implications of this shift, which began in the late 1980s and accelerated throughout 1990, are enormous for users and suppliers alike. For users, the availability of supercomputing power at desktop system prices means far greater access to such power for scientists, engineers and other technical users.

For vendors, the change from host-based supercomputing to a distributed model threatens current market leaders and gives hope and opportunity to others. Dataquest Inc., a San Jose, Calif.-based market research firm, estimates worldwide demand for technical-computing systems will grow a vigorous 14% per year from \$22.6 billion in 1989 to

\$43.5 billion in 1994.

The full impact of this change in technical computing won't, of course, be known for some time. What is known today is information about the technologies, products and suppliers available to users in search of economical solutions to the technical challenges they face.

For our purposes, these alternative supercomputing machines will be known as "supersubstitutes," and will include minisupercomputers, graphics supercomputers (or superworkstations) and specialized parallel computers. Workstations, superminicomputers and mainframes can be considered less optimal supersubstitutes.

Benchmarking Super Performance

Cost-effective FLOPS—that is, the power to process large amounts of floating point operations per second—come in many new forms. Machines as varied as a workstation from Sun Microsystems Inc. or a special-purpose computer from Thinking Machines Corp. all do the same computation as a supercomputer—for less than 5 to 50% of the price. Evidence of this capability abounds.

Take the 1989 results of the PERFECT benchmark, for example. Performance Evaluation for Cost-Effective Transformation, a suite of tests developed at the University of Illinois's Center for Supercomputing Research and Development in Urbana, attempts to measure supercomputer performance and cost-effectiveness.

In the 1989 trials, an eight-

processor Cray Y-MP/832 supercomputer took the laurels for peak performance by achieving 22 million floating point operations per second for unoptimized baseline programs and 120MFLOPS for hand-tuned, highly optimized programs. In contrast, a single-processor Stardent 3000 superworkstation performed at 4.2MFLOPS with no tuning. Slower, but 27 times more cost-effective than the untuned Cray Y-MP programs, considering the price difference. The Y-MP sells for \$18.8 million, while the Stardent lists for \$123,000.

MINISUPERCOMPUTERS

Strengths: Have computational capacity approaching one Cray Y-MP processor. Support a modestly interactive, distributed mode of use. Exploit the gap left when Digital began in earnest to ignore its technical user base. Price/performance is roughly the same as a supercomputer.

Weakness: Overpriced and underpowered when compared with newer classes of technical computers.

Examples: Machines from Alliant Computer Systems Corp. and Convex Computer Corp.

Advice: Watch how suppliers respond to pressure from workstation and superworkstation rivals. Alliant plans to improve price/performance by migrating to Intel's i860 RISC chip as a source of cost-effective MFLOPS. Convex is likely to follow the Cray path of achieving higher clock speeds by using emitter-coupled logic (ECL) technology.

OPTION
1



BELL: Superworkstations are a powerful alternative to big iron for users of traditional supercomputers.

The PERFECT results typify the inverse relationship between size of computer and computing power evident throughout the worldwide computer industry as desktop systems (PCs, workstations, etc.) are compared with host models (mainframes, minicomputers, etc.). When it comes to getting high-performance computation for scientific and engineering problems, the biggest machines rarely are the most cost-effective.

Even the world's most influential community of supercomputer users—the researchers who use, at no cost, the National Science Foundation's six (and soon to be five) supercomputing centers in the United States—get relatively little processing power per year from the big machines at their disposal. About 30 Cray X-MP processors are available to these researchers, which represent about 240,000 processor hours per year. But there are so many researchers (10,000) that each gets little more than 24 hours per year of supercomputer use. That's just what a high-power PC can deliver in a month.

There are many other factors besides machine availability that inhibit researchers from getting the most out of supercomputers, including the investment they must make in understanding

SUPERWORKSTATIONS

Strength: Provide between 10 and 33% of the computational capacity of a single-processor Cray Y-MP at less than 3% of the cost.

OPTION 2

Weakness: Peak benchmark speed is only 2% of eight-processor Cray Y-MP.

Examples: These machines combine varying degrees of supercomputer capacity with integral 3-D graphics capabilities for project and departmental use (i.e., multiple users per system) at costs ranging between \$50,000 and \$200,000. Superworkstations from IBM, Hewlett-Packard, Silicon Graphics and Stardent—priced aggressively at \$25,000 to \$50,000 per machine—make similar features affordable for personal use. Notable is IBM's RS/6000 superscalar workstation, on which several researchers have run programs at the same speed as on a Cray Y-MP.

Advice: Distributed approach makes superworkstations more effective than supers.

vector processing, parallel processing, memory hierarchies and other elements crucial to effective use of supercomputers.

Large, central supercomputing facilities are not necessarily flawed as a source of compute power, unless they attempt to be a one-stop solution. They may be the best resource for the very largest us-

□ SUPERWORKSTATIONS HOLD THE MOST PROMISE BECAUSE THEY EMBODY MORE LEADING-EDGE TECHNOLOGIES.

ers with big, highly tuned parallel programs that require any or all of the following: large memories, file capacity of hundreds of gigabytes, archival files and the capacity to share large databases and large programs. Large supercomputers also suffice for the occasional user who needs only a few hours of computing a year and doesn't want to own or operate a computer.

MASSIVELY PARALLEL COMPUTERS

Strength: Can supply same peak delivered power (1.5GFLOPS) on select applications as a Cray Y-MP/8.

OPTION 3

Weakness: Not suitable for general scientific work loads such as chemical modeling, whose applications can't readily take advantage of massively parallel architectures.

Examples: Machines come in two forms. Multicomputers, which connect hundreds of microprocessors together, are made by the likes of Intel, Inmos Corp. and NCube. Single-instruction, multiple datastream (SIMD) machines come from Thinking Machines, MasPar and others. The notable exception is newcomer Kendall Square Research Corp.'s multiprocessor computer—scalable up to 1,000 processors. It operates equally well for both massive transaction processing and massively parallel computation. Unlike the SIMD or multicomputers, the Kendall Square multiprocessor is a general-purpose computer.

Advice: Multicomputers and SIMD are not directly substitutable for current supers. With time, compilers should be able to better exploit these architectures.

But the biggest iron is not very well suited to the needs of the majority of users working on engineering or scientific problems. The big supercomputers lack the interactive and visualization capabilities that computer-aided design requires, for example. As a result, even with free computer time, only a small fraction of the research community—5 to 10%—uses the NSF centers. The rest are buying smaller computing resources instead. To such researchers, less is indeed more.

Shopping for Super Substitutes

A user can opt for a super substitute if the system supplies at least 10% of the power of a conventional supercomputer. That's enough muscle for a supersubstitute to deliver at least the same amount of computation in one day that the typical user could expect from a large, time-shared supercomputer in one-half to two hours. The user should use three other criteria in selecting a super substitute: price/performance leadership in its class, high throughput on a wide variety of jobs and appropriate memory and other resources.

Numerous machine types are contenders as supersubstitutes. Here's a rundown of them:

- **Workstations, minicomputers and superminicomputers.** They are the most common technical computing platforms today.
- **Mainframe computers.** Those with vector processing features also rival supers in some cases.
- **Massively parallel systems.** Such machines contain hundreds, even thousands of microprocessors capable of working simultaneously on a single program. But

MAINFRAMES

Strength: Can provide up to 25% of peak Cray Y-MP's power at as little as 50% of the price.

OPTION 4

Weakness: Many mainframes cost more than some of the low-end Cray models. The most representative machine, IBM 3090/600 with vector facilities, has only one-third of the capability of Cray Y-MP/8 to carry out a given work load.

Examples: IBM 3090s, Digital VAX 9000s.

Advice: No practical cost benefit from substituting one centralized, time-shared resource for another.

SYSTEMS

TECHNICAL COMPUTING

many supercomputer applications have to be completely reprogrammed—an expensive investment by a user—to exploit the advantage of parallel architectures.

■ **Minisupercomputers.** Minisupers, as they are known, were the first viable supersubstitutes, appearing in 1983. And they were a cost-effective source of supercomputing capacity until 1988, when superworkstations first came on the scene.

■ **Superworkstations.** This machine class is the most vigorous of all technical computer categories, attracting the ma-

□ BY THE YEAR 2000, NEARLY ALL PCs WILL HAVE THE CAPABILITY OF TODAY'S SUPERCOMPUTERS.

majority of buyers and supplying the bulk of the capacity for high-performance technical computing.

Of all these contenders, superworkstations hold the most promise for the decade ahead because they embody more leading-edge technologies than any other class of technical computer. Moreover, the benefits of these technologies—such as reduced instruction set computing (RISC) and complementary metal oxide semiconductors (CMOS)—are readily apparent to users. Not only do the superworkstations deliver improved price/performance and better compiler techniques than supercomputers—standard fare for the leading edge—they also enable users, via better interfaces, to more easily interact with programs.

WORKSTATIONS

Strength: Provide up to 10% of the capacity of a Cray Y-MP processor at less than 2% of the cost.

**OPTION
5**

Weakness: Relatively slow speed prevents them from hitting performance peaks in compute-intensive programs for which the vector and parallel capabilities of supercomputers were developed.

Examples: Machines from Digital, Hewlett-Packard, Silicon Graphics and Sun.

Advice: Not ideal as super substitutes.

Supersurvey of the Decade

What are your technical computing plans for the 1990s?

	Supers	Mini-supers	Main-frames	Mini/Super-minis	Super work-stations	PCs
What's your chief technical computing platform today?	325	326	327	328	329	330
What will be your chief technical computing platform in 1995?	331	332	333	334	335	336
What will be your chief technical computing platform in 2000?	337	338	339	340	341	342

Which technical computing suppliers do you want to hear from immediately?

Alliant	343	Kendal Square Research	351
Convex	344	MasPar	352
Cray Research Inc.	345	NCube	353
Cray Computer Corp.	346	NEC	354
Digital	347	Silicon Graphics	355
HP	348	Stardent	356
IBM	349	Sun	357
Intel	350	Thinking Machines	358

What applications do you have that require extraordinary computational or graphical performance?

General-purpose computing	359
Animation	360
Image or signal processing	361
Mechanical CAD/CAM	362
Electrical CAD/CAM	363
Chemical modeling or computational chemistry	364
Computational fluid dynamics	365
Visualization of results from another computer	366

Source: DATAMATION; Authors.

Radically new applications will spring up around superworkstations in much the same way that new uses for systems developed around the personal computer. And, like their PC counterparts, these applications won't just be dressed-up versions of the code museums that populate the supercomputer, mainframe and minicomputer worlds. Although the scientific and engineering fields will generate most of the radical new ideas, expect contributions from the commercial sector, as well. Several financial institutions, for example, are applying superworkstations for econometric modeling, work flow optimization or portfolio analysis.

Because these machines are based on fast-evolving technologies, their performance can be expected to improve by 50% annually during the next five years. The speed with which they execute programs also will accelerate, with clock rates of more than 100 megahertz possible by 1992.

What's most promising, perhaps, about the future of superworkstations is their natural appeal to user organizations already prone to invest in distributed-computing technologies. U.S. corporations and their European peers are much quicker to embrace distributed computing than to invest in supercomputers, mainframe systems and minisupers—big ticket items that can be painstakingly slow and convoluted to purchase and that have prices that limit their acceptance and application. Superworkstations, by their distributed nature, have a better chance of attracting creative, productive users with authority to purchase machines in the \$10,000-to-\$50,000 range.

If, in fact, superworkstations assume such a significant role in the 1990s, what then will become of the supercomputer? For now, the supercomputer continues to be a protected species because of its use in U.S. Department of Defense applications. Like the Harley-Davidson motorcycle, the U.S. supercomputer has be-

come a token symbol of trade and competitiveness. Japanese manufacturers have begun to make computers with peak speeds equal to or greater than those from Cray Research Inc. or Cray Computer Corp.

From an economic standpoint, U.S. and European information technology suppliers are fortunate that their Japanese rivals are expending the large development resources that supercomputers

such as NEC Corp.'s SX-3 require; otherwise engineers in Tokyo and Osaka could be developing consumer and workstation products that could have a far greater impact on Western markets.

In reality, the supercomputer is becoming a special purpose computer that

MINISUPERS, AS THEY ARE KNOWN, WERE THE FIRST VIABLE SUPERSUBSTITUTES.

is really only cost effective for highly parallel problems. It has about the same performance as highly specialized, parallel computers like the Thinking Machines' Connection Machine, NCube's systems of the same name and Intel Corp.'s iPSC products. Yet the supercomputer costs 10 times as much as any of those parallel systems because of its expensive circuit and memory technology.

The evolution of the traditional supercomputer must change to a more massively parallel and scalable structure if it is to keep up with the peak performance

of evolving new machines. By 1995, specialized massively parallel computers capable of a trillion floating point operations per second will be available to simulate a much wider range of physical phenomena.

By the year 2000, nearly all PCs will have the capability of today's supercomputers. Such functionality will enable users to simulate immense and varied models—the very basis of technical computing. In the intervening years, superworkstations look to be the technology that engineers and scientists will harness to achieve the design, medical and other breakthroughs that will distinguish their professions.

Gordon Bell is chief scientist at Stardent Computer Co., a superworkstation manufacturer in Sunnyvale, Calif. He formerly worked for Digital Equipment Corp., where he was responsible for the company's VAX family of midrange systems and workstations. Glen Miranker was the chief architect for Stardent's 3000 product line. He previously worked for Valid Logic Systems Inc. and for IBM.

MINICOMPUTERS AND SUPERMINIS

Strength: Provide up to 20% of the capacity of a Cray Y-MP processor at less than 16% of the cost.

OPTION 6

Weakness: Only 0.25% of the peak speed of a Cray Y-MP.

Examples: Midrange machines from Digital, HP and others.

Advice: Less than ideal super substitutes.

DID YOU KNOW...

Twice each month

DATAMATION's Career Opportunities Section

brings you the latest in current Job Opportunities in the MIS market, where you can reach

- Software Engineers • Systems Integrators
- Data Processing Managers • Programmers/Analysts.

See Career Opportunities beginning on Page 127