Minisupercomputers —Advancing Technologies

Minisupercomputers have become something of an industry phenomenon during the last few years. They have been marketed as "Crayettes," "gap fillers," "VAX-killers," and "VAX-alikes." The appellations reveal much about the way the machines are perceived and sold. Minisupers fill a pronounced performance gap between Digital Equipment Corporation supermini VAXs at the lower end and full-blown Cray Research supercomputers at the high end. In fact, many minisuper product lines incorporate features from both VAX superminis and Cray supercomputers to encourage user migration to their systems. At least one product uses the Cray X/MP instruction set and others use Cray Fortran extensions, leading some industry observers to dub them Crayettes. Others use VAX Fortran extensions and emulate VAX/VMS environments.

For virtually all the new minisuper products, the marketing strategy is the same; they are designed to provide users with near-supercomputer performance at the price of a supermini. Minisupers sell for between $100,000 for basic entry-level systems to more than $1 million for systems that overlap the performance of a low-end supercomputer.

To achieve near-supercomputer performance at a supermini price, minisuper start-ups and established industry players have developed products using mostly inexpensive, off-the-shelf microprocessor technology. To make common microprocessor components run like thoroughbreds, vendors have come up with innovative architectures combined with intelligent compiler technology. This approach is not typically available to conventional, general-purpose superminis. Most vendors resort to parallel processing techniques to attain better performance. To handle numeric-intensive arithmetic operations, system designers typically combine parallel processing with specialized hardware to handle vectors and floating-point operations.

Within the last year, the minisupercomputer has gone from a curiosity to a hot new product. More start-ups have entered the market, and established vendors have already announced second-generation systems. With so many new players and new products chasing the same base of technical users, price-cutting and confusion seemed inevitable. If Digital Equipment Corporation and IBM enter the market, as they are expected to do in the near future, the high-performance system fray should get even uglier.

To make sense of these events, Datapro introduces a new report on the minisuper. The report examines the hardware and software technologies and takes a closer look at six major players. Their offerings are explained through detailed vendor profiles, tables, and diagrams.

Other vendors have achieved near-supercomputer performance largely through software. Multiflow Computers, for instance, uses a single processor built to process very long instruction words (VLIWs) up to 1024 bits in length. Another start-up, Cydrome Inc., uses a dataflow concept, an approach similar in some respects to the VLIW technique. These approaches reformat existing code and process larger portions of an application simultaneously, using fine-grained parallelism. The vendors claim their approaches are superior to vector and traditional multiprocessor parallel processing, since a greater portion of code can be optimized. Despite their limitations, machines using parallel and vector architectures continue to be more popular, simply because their technology has been commercially available the longest.

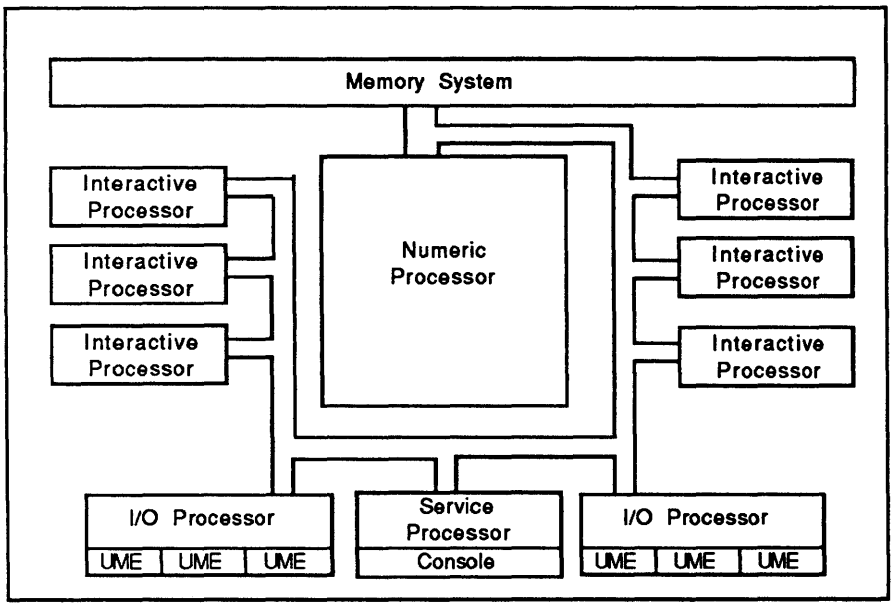

Figure 1. The Cydra 5 system diagram shows the five basic system components used to implement the company's dataflow architecture. The Numeric Processor contains multiple functional units that perform various operations. Up to 56 parallel operations can be under way within the numeric processor and up to 7 new operations can be initiated in each machine cycle.

Minisupercomputers —Advancing Technologies

TABLE 1. MINISUPERCOMPUTERS AT A GLANCE

MODEL	Alliant FX/4	Alliant FX/40, FX/80	Convex C120	Convex C130	Convex C210-C240
SYSTEM CHARACTERISTICS					
Date announced	October 1987	February 1, 1988	March 1988	March 1988	March 1988
Date first delivered	October 1987	First-quarter 1988	March 1988	March 1988	C210, Mar. 1988; C220, 2nd-qtr. 1988; C230-C240, 4th-qtr. 1988
Number of processors	1 to 4	FX/40, 1 to 4; FX/80, 1 to 8	1	1	1 to 4
CPU cycle time, nanoseconds	—	—	100	55	40
Cooling technique	Air-cooled	Air-cooled	Air-cooled	Air-cooled	Air-cooled
Operating systems	Concentrix (UNIX 4.2)	Concentrix (UNIX 4.2)	Convex UNIX 4.2 BSD	Convex UNIX 4.2 BSD	Convex UNIX 4.2 BSD
SOFTWARE CHARACTERISTICS					
Languages supported	Fortran, C, Ada, Assembler, Portable C, Pascal, Lisp, and APL	Fortran, C, Ada, Assembler, Portable C, Pascal, Lisp, and APL	Fortran, C, Ada	Fortran, C, Ada	Fortran, C, Ada
Name of optimizing compiler(s)	FX/Fortran, FX/C, FX/Ada, Pascal	FX/Fortran, FX/C, FX/Ada, Pascal	Convex Fortran, Convex C	Convex Fortran, Convex C	Convex Fortran, Convex C
MEMORY CHARACTERISTICS					
Memory range, bytes	Up to 128M	FX/40, up to 128M; FX/80, up to 256M	32M to 1G	Up to 1G	Up to 2G
Memory bandwidth	188.8M bytes/sec	188.8M bytes/sec	80M bytes/sec	145M bytes/sec	200M bytes/sec
Memory interleaving	4 way per memory module	4 way per memory module	8 way to 32 way	32 way	Up to 64 way
Memory type	256K-bit dynamic RAM	256K-bit dynamic RAM	—	—	—
I/O CHARACTERISTICS					
Number of I/O processors	1 to 6	FX/40, 1 to 6; FX/80, 1 to 12	5 channel control units (CCUs)	4 CCUs	4 to 8 CCUs
Channel speed	5M bytes/sec	—	4M to 80M bytes/sec	4M to 80M bytes/sec	4M to 80M bytes/sec
COMMUNICATIONS PROTOCOLS	Ethernet, TCP/IP, DECnet, NFS, NCS, X-Window System, NeWS, X.25, DDN, HYPERchannel	Ethernet, TCP/IP, DECnet, NFS, NCS, X-Window System, NeWS, X.25, DDN, HYPERchannel	Ethernet, TCP/IP, NFS, COVUEnet, HYPERchannel	Ethernet, TCP/IP, NFS, COVUEnet, HYPERchannel	Ethernet, TCP/IP, NFS, COVUEnet, HYPERchannel

A dash (—) indicates information was not available.

Customers purchasing minisupers are using them to process many of the same applications they now run on a supermini. Major applications include aerodynamics modeling, structural analysis, seismic processing, reservoir modeling, finite element modeling, circuit simulation, image processing, and artificial analysis. Targeted markets include government-related installations; aerospace; petroleum/geophysics; and highly commercial segments, such as auto manufacturing, computational chemistry, CAD/CAM/CAE, econometric modeling, graphics, and film animation.

Many users would rather run these applications on a full-scale supercomputer, but since supercomputers are financially out of their reach, many departmental users will gladly settle for a minisuper. A minisuper provides users with better single-job turnaround than is available on a supermini.

Similar to supercomputers, minisupers have been a boon to more users in the commercial sector. Such systems let users simulate and test prototypes or simulate real-life situations. Designers can test engine designs, car bodies, or airplane aerodynamics and make design changes using computer-aided design approaches. Before developing computer-simulated models, aerospace and auto manufacturing firms were forced to build expensive prototypes and equally expensive wind tunnels to test and analyze prototypes. Of course, industry continues to use wind tunnels, but minisupers now provide users with cost-effective alternatives.

High-performance systems have been a particular boon to the oil industry, which uses such systems to analyze seismic data and for reservoir modeling. The minisuper uses the data to reconstruct the underground strata of a proposed oil drilling site. Engineers can then analyze the computer findings and pinpoint where oil is likely to be and, more important, where it's not likely to be. This could then eliminate the expense of drilling a dry hole. The resulting savings in exploration costs more than justifies the cost of the vector hardware.

Who Makes What?

Minisuper players are an interesting collection of start-ups and a few established minicomputer vendors that have extended their product lines to meet the need of scientific users who require more power. Early players include Alliant Computer Systems and Convex Computer Corporation, start-ups that entered the market in 1985, and Scientific Computer Systems (SCS) Corporation, which entered the market in 1986. While SCS got off to a slow start, Convex and Alliant have now emerged as clear minisuper market leaders. Convex, which delivered its first minisuper in 1985, now claims to have the single largest minisuper installed base. The company has shipped more

Minisupercomputers —Advancing Technologies

TABLE 1. MINISUPERCOMPUTERS AT A GLANCE (Continued)

MODEL	Floating Point M64 Series	Multiflow TRACE 7/100, 7/200, 14/200	Prime/Cydrome Cydra 5	Scientific Computer Corporation-30/XM	Scientific Computer Corporation-40/XM
SYSTEM CHARACTERISTICS					
Date announced	1986	7/200, April 1987; 7/100 and 14/200, February 1988	January 1988	January 1988	—
Date first delivered	1986	7/200, July 1988; 7/100, February 1988; 14/200, May 1988	January 1988	Second-quarter 1988	July 1986
Number of processors	M64/20-M64/60, 1; M64/145, 3 to 31	1	1	1	1
CPU cycle time, nanoseconds	—	7/100, 170	40	—	45
Cooling technique	—	Air-cooled	Air-cooled	Air-cooled	Air-cooled
Operating systems	System Job Executive	TRACE/UNIX (Unix 4.3 BSD)	Cydrix 5.3 (UNIX System V.3)	CTSS, COS, SCENIX (UNIX System V.3)	CTSS, COS, SCENIX (UNIX System V.3)
SOFTWARE CHARACTERISTICS					
Languages supported	Fortran, C, Assembler language	Fortran, C	Fortran 77 with Digital Equipment, IBM, UNIX, and Cydrix extensions	LRLTRAN (extended version of Fortran)	LRLTRAN (extended version of Fortran)
Name of optimizing compiler(s)	APFTN64 (Fortran 77); APAL64 (Assembler); APC64 (C)	Fortran with VAX/VMS, -8X, and IBM extensions	Cydrix Fortran 77	CFT, Civic (Fortran based)	CFT, Civic (Fortran based)
MEMORY CHARACTERISTICS					
Memory range, bytes	M64/40, 50, and 60, 1 to 9 megawords	16M to 512M	8M to 256M	1, 2, or 4 megawords, expandable to 256 megawords	1, 2, or 4 megawords, expandable to 256 megawords
Memory bandwidth	—	490M bytes/sec	400M bytes/sec	—	711.1M bytes/sec
Memory interleaving	—	64 way	Up to 64 way	—	16 way
Memory type	M64/40, dynamic RAM; M64/50 and 60, static RAM	256K-bit and 1M-bit dynamic RAM	256K-bit and 1M-bit chips	—	256K-bit NMOS
I/O CHARACTERISTICS					
Number of I/O processors	M64/40, 50, and 60, 1 to 7	—	2	—	2 to 8
Channel speed	M64/40, 50, and 60, 38M to 44M bytes/sec	7/100, 94M bytes/sec	40M bytes/sec	—	44.5M bytes/sec per I/O module
COMMUNICATIONS PROTOCOLS	—	Ethernet, TCP/IP, NCS, NFS, NeWS, X Windows, DECnet, Hasp	Ethernet with TCP/IP	—	—

A dash (—) indicates information was not available.

than 260 machines to more than 150 customers in 27 countries worldwide. Floating Point Systems (FPS), a firm that predates the minisuper era, disputes Convex' market leadership claims. FPS notes it has been selling 64-bit processor attachments since 1981, making the company a minisuper pioneer. FPS has installed more than 400 systems that it classifies as minisupers. Although FPS systems require a supermini or mainframe host, FPS contends its systems should be considered minisupers since they compete against other minisuper vendors. Additionally, FPS notes standalone systems from Convex and Alliant don't really stand alone in many instances, since they also can be used as compute-intensive servers in a supermini or scientific workstation network.

Convex, Alliant, SCS, and FPS were part of the first industry wave and have developed a degree of name recognition. They have since been joined within the last couple of years by at least a half-dozen new companies. The most notable newcomers among the bunch include Multiflow of Branford, Connecticut, Cydrome of Milpitas, California, and Gould of Fort Lauderdale, Florida. Of these players, Cydrome recently joined forces with Prime to market its minisupers. Meanwhile, Gould, an established supermini vendor, expanded into the minisuper market in 1987.

The emergence of so many new players in so short a time has led to cutthroat competition and heavy discounting. Many of the established vendors such as Convex, Alliant, and Floating Point are finding it harder to close new sales now that so many new products compete in the same price/performance range. As expected, many analysts are already predicting an industry shakeout. Culler Scientific Systems, which introduced its product line in 1986, went under a year later. Floating Point Systems, the best known of the established players, reported big losses in 1987 and was forced to trim staff and manufacturing facilities.

Despite the dire warnings, numeric-intensive computing continues to be a growth market for the time being. By contrast, traditional mainframe and mini markets remain relatively flat. Analysts who study the market generally like to divide it into three distinct but overlapping segments. Supercomputers and high-performance mainframes populate the high end. While these systems are coming down slightly in price as more competitors enter

Minisupercomputers —Advancing Technologies

TABLE 2. MINISUPERCOMPUTERS: PRICE/PERFORMANCE

MODEL	Alliant FX/4	Alliant FX/40, FX/80	Convex C120	Convex C130	Convex C210-C240
PERFORMANCE					
MIPS	—	—	11M Whetstones*	20M Whetstones*	31M to 124M Whetstones*
FLOPS: peak performance	47.2M*	94.4M to 188.8M*	20M*	35M*	50M to 200M*
FLOPS: Linpack model	—	30M to 65M, fully configured**	—	—	—
FLOPS: Livermore Loop	—	—	—	—	—
PRICE RANGE	$99,900	$149,000 to $299,000	$275,000	$445,000	$595,000 to $1.35 million

**Vendor supplied data.*
***Argonne National Laboratory, January 12, 1988; the benchmark measured how fast tested machines solved a system of linear equations in full precision.*
A dash (—) indicates information was not available.

the field, supercomputers continue to be the most expensive and most powerful commercially available systems in the world. Minisupers and parallel processors dominate the mid-range performance spectrum. As noted, in power terms, they're positioned between a high-end supermini and a low-end supercomputer. High-performance scientific workstations, many with sophisticated 2-D and 3-D graphics capabilities, round out the low end.

International Data Corporation (IDC), the Framingham, Massachusetts market research firm, predicts that the value of worldwide supercomputer and high-performance mainframe shipments is expected to grow at a compound annual rate of 25 percent between now and 1991. By comparison, IDC projects the value of all large-scale systems including traditional mainframes will grow only 3 percent; and minisupers and parallel processors are expected to have a combined 1986 through 1991 compound annual growth rate of 46 percent, compared to 7 percent for the entire medium-scale market. By 1991, IDC predicts high-performance systems should account for about 15.0 percent of the large-scale market, 5.5 percent of the medium-scale market, and 7.0 percent of the total multisystem market. Overall, IDC projects the high-performance systems sector will grow from the estimated $1.15 billion market of 1986 to a $4.16 billion market by 1991.

Although different market segments will grow at different rates, each market addresses many of the same technical users. These markets tend to overlap, since many technical users access all three levels of scientific computing through sophisticated networking setups. Most minisuper vendors, for instance, provide industry-standard communications protocols and gateways allowing users at different hardware levels to link their systems to larger or smaller systems. Minisuper users can use a workstation, for instance, to graphically display the results of a job processed on the minisuper. Additionally, workstation users can run numerically intensive portions of an application on a minisuper to attain greater performance or turnaround time. Recognizing the complementary nature of workstation and minisuper product lines, several leading workstation vendors, including Sun, Apollo, and Silicon Graphics, have entered into joint marketing agreements with minisuper vendors. Under these arrangements, workstation and minisuper vendors both agree to market each other's products, particularly when they call on customers looking for both.

The Secret of Their Success

Start-ups that have already established a track record in the minisuper market owe much of their success to Digital Equipment. The name Digital has become synonymous with scientific computing at the mini and supermini levels. In recent hears, however, the Maynard, Massachusetts company has directed more of its attention to commercial processing. Many believe Digital's preoccupation with commercial processing, together with its campaign to steal market share away from IBM, could be causing the company to neglect its traditional base of engineering/scientific users. Many VAX uses have already installed Convex, Alliant, and other competing minisupers.

Of course, Digital, which has come to dominate the mid-range marketplace with the installation of thousands of VAXs, could easily overtake its minisuper rivals with a minisuper of its own. A Digital Equipment minisuper equipped with vector processors would find a ready market among many existing VAX customers looking for more power. Additionally, such a product would close up the well-publicized performance gap between a high-end VAX and a Cray. At the moment, Digital is believed to be working on a minisuper, which industry insiders believe may be ready by 1989. Software development problems may be holding up the announcement, according to industry speculation. How the Digital minisuper will fit into the VAX marketing mix could be the source of an even larger problem. A minisuper with better price/performance than less powerful VAXs could hurt sales in other VAX performance ranges.

In an ironic turn of events, Digital's chief commercial rival, IBM, has taken a much bigger interest in engineering/scientific computing in recent years. IBM plans to introduce a product for the mid-range technical computing market, perhaps by 1989 or 1990. The company already offers the Vector Facility (VF) attachment for its 3090 mainframe users. IBM could be planning to offer the VF to its mid-range customers or could be planning to introduce a brand-new product. As usual, IBM isn't talking specifics.

Minisupercomputers —Advancing Technologies

TABLE 2. MINISUPERCOMPUTERS: PRICE/PERFORMANCE (Continued)

MODEL	Floating Point M/64 Series	Multiflow TRACE 7/100, 7/200, 14/200	Prime/Cydrome Cydra 5	Scientific Computer Corporation-30/XM	Scientific Computer Corporation-40/XM
PERFORMANCE					
MIPS	—	7/100, 41; 7/200, 53; 14/200, 107*	—	16*	22*
FLOPS: peak performance	6M to 341M*	7/100, 11M; 7/200, 15M; 14/200, 30M*	—	33M*	44M*
FLOPS: Linpack model	M64/60, 5.6M; M64/30, 2.1M**; M64/145, 24M to 101M*	7/100, 4.2M; 7/200, 6M; 14/200, 10M*	10M**	—	8M**
FLOPS: Livermore Loop	—	7/100, 1.8M; 7/200, 2.3M; 14/200, 3.4M*	3.7M*	—	—
PRICE RANGE	M64/145, $395,000 to $1.75 million	$197,500 (7/100) to $399,500 (14/200)	$549,000, Model 1200 configuration	$371,000	$595,000

**Vendor supplied data.*
***Argonne National Laboratory, January 12, 1988; the benchmark measured how fast tested machines solved a system of linear equations in full precision.*
A dash (—) indicates information was not available.

Digital's apparent indecisiveness in the minisuper segment has created opportunities for new minisuper rivals openly courting the VAX user. Most offer software products to implement VAX/VMS environments on their respective minisupers, along with VMS Fortran conversion tools, Fortran compilers with VMS extensions, and VMS Command Language shells. VAX-related products are designed to make the migration from a VAX and to a non-VAX system as painless and transparent as possible. In some instances, vendors claim VAX users can recompile their VAX Fortran code on VAX-compatible systems with little or no programmer intervention. Likewise, VMS-like operating system shells provide former VAX'ers with a familiar environment that eliminates the need to learn new commands and syntaxes.

To lure VAX users to minisupers, vendors constantly publish relative performance ratings. Many compare their systems directly with a VAX 8700. Convex, for instance, claims its six new systems announced in March achieve up to five times the performance of a VAX 8700 and overlap the performance of a Cray X-MP/14. Alliant claims its second-generation systems announced in February deliver up to 15 times the performance of a VAX 8700 in certain computationally intense applications.

How Minisupers Compare

To compare their various systems to other minisupers and full-blown supercomputers, vendors turn to industry-standard performance benchmarks. Generally, vendors marketing high-performance systems measure performance in floating-point operations per second (FLOPS), which is the number of floating-point operations that can be performed by the system's central processing unit in one second. Most vendors release data that describes the peak performance of their various systems. Virtually all minisupers perform more than 1 million FLOPS, expressed as *mega*flops. Within the minisuper segment, peak performance generally ranges from 10 megaflops at the low end to 200 megaflops at the high end. Systems that perform in the 200-megaflop range or better overlap the performance of full-blown supercomputers. Today's most powerful supers can now operate at a peak performance that exceeds 1 billion FLOPS, expressed as *giga*flops. Supercomputer peak performance starts at about 100 megaflops at the low end and approaches 10 gigaflops at the extreme high end.

Peak performance, at best a theoretical measurement of raw power, is generally used to help users and potential customers compare competing systems and make product evaluations. Vendors concede, however, that users will rarely achieve peak performance running real jobs. Because of hardware limitations and software constraints, processors and functional units in many super- and minisupercomputers cannot be kept uniformly busy to sustain peak performance.

Performance in the real world depends on the type of problems users intend to solve, the percentage of code that can be reformatted to run more efficiently on the specialized supercomputer hardware, and the degree of performance degradation attributable to memory access and input/output bottlenecks. To provide a more realistic relative performance benchmark, vendors also release the results of Linpack and Livermore Loop tests, two widely recognized benchmarks for testing specific supercomputing performance. The Linpack benchmarks, used by the Argonne National Laboratory, compare relative performance for a range of computer processors from micros, minis, superminis, and minisupers to full-scale supercomputers. Linpack tests the capability of a processor to solve multiple linear equations. The Livermore Loops are a series of benchmark programs developed by the Lawrence Livermore Laboratory for measuring floating-point computational performance. It tests the computer's capability to solve more than one multiple CPU-intensive equation.

Some routinely published industry-standard benchmarks show many minisupers achieve performance on certain applications that's competitive with more expensive low-end supercomputers. Minisupers achieve this performance using vector and parallel architectures that at one time were only available on more expensive supercomputers. Other minisuper architectures grew out of experimental university-based research projects. Vendors selling

Minisupercomputers —Advancing Technologies

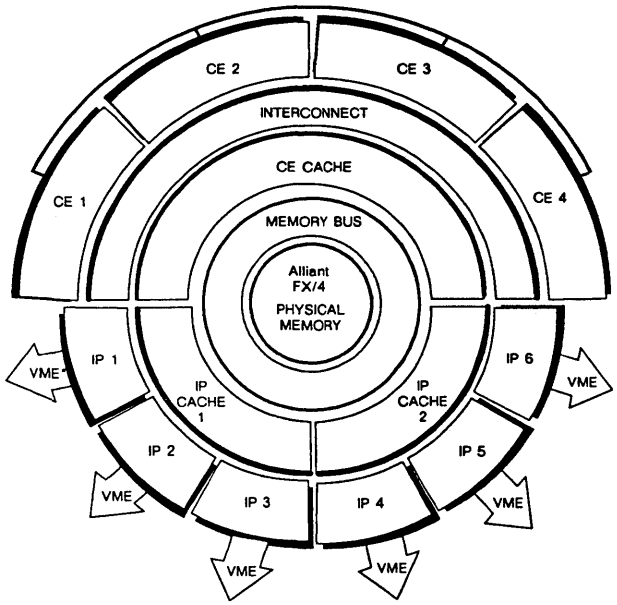

Figure 2. Alliant FX/4 system features Computational Elements (CEs) for handling compute-intensive vector and scalar code and Interactive Processors (IPs) for running interactive user jobs and for executing the operating system and I/O tasks. For maximum throughput, multiprocessing applies up to four CEs and six IPs to the execution of independent jobs.

VLIW and dataflow-based products now claim their system designs have overcome technical problems, making the approaches commercially practical for perhaps the first time.

Minisupers, supercomputers, and high-performance mainframes equipped with specialized hardware for numeric-intensive computing all have something in common. Most of the machines are designed to process multiple datastreams and instruction streams simultaneously, using some form of parallelism. Vendors have had to turn to parallel processing techniques at both the minisuper- and supercomputer levels largely because the gains achieved through single-processor architectures alone are almost exhausted. Some systems achieve parallelism using multiple processors, while other systems achieve this using one processor incorporating multiple functional units, instruction pipelining, or other forms of internal parallelism. Multiflow claims to achieve parallelism mostly through software.

To make it easier to categorize some of these architectures, vendors often turn to a popular classifying scheme first developed in the 1960s by Michael J. Flynn. The Flynn approach groups machines by how they handle both streams of instructions and streams of data. The four categories of machine organization are:

- Single instruction stream-single datastream (SISD) machines, which encompass most uniprocessors in use today. These machines process data sequentially, although enhancements, such as pipelining and memory interleaving, may also be used.
- Multiple instruction stream-multiple datastream (MIMD) machines, which employ multiple processors to handle multiple instruction streams. Multiprocessor configurations and parallel processors fall into this category.
- Single instruction stream-multiple datastream (SIMD) machines, which may employ vector and array processing. This architecture is often used for scientific work involving data arrays featuring fixed intervals between successive elements. Array machines, for instance, will typically use multiple processing elements to execute a single instruction on arrays of floating-point add and multiply data simultaneously.
- Multiple instruction stream-single datastream (MISD) machines, which employ two or more processors that perform separate instructions on the same data. Since this approach is thought to be impractical, it has received little attention over the years.

What Makes Minisupers Super?

To achieve high performance in numeric-intensive applications, most supercomputers, minisupers, and high-performance mainframes employ specialized arithmetic units to perform vector processing. Integrated vector processing hardware is ideal for speeding up the processing of repetitive data expressed in the form of arrays or tables, a type of processing often occurring in scientific work. Vector processing reformats data into arrays, making it possible to carry out repetitive steps or perform multiple operations using one instruction. Vector processing is especially useful for Fortran applications that involve looped instructions, a multiple-step, logical operation that will bog down scalar machines. Many scientific problems are so big, they would be otherwise unsolvable without a vector processor. Generally, mainframe and most commercial processors use scalar processing approaches, which carry out operations sequentially, one operation at a time. Using the scalar approach, it can take many days to calculate the same data that a vector processor could tackle in a matter of minutes.

Most vector machines make extensive use of pipelines that overlap key functional operations to achieve high processing speeds. Machines can use pipelines to perform instruction and operand fetch for both scalar- and vector-type operands. Special arithmetic pipelines can be dedicated to both scalar and vector operations.

Vector machines also use relatively large memories and high-speed caches to handle larger portions of a program at one time and cut down on time-consuming I/O fetches. To reduce memory access conflicts, such machines use a high degree of memory interleaving. Supercomputers tend to use a higher degree of interleaving than minisupers. Minisupers may feature 32-to-64 way interleaving, while a supercomputer may feature 128-way and 256-way interleaving.

Minisupercomputers —Advancing Technologies

Generally, vector systems will divide a task up among a vector processor and a very high speed scalar processor. Since not all operations will be vectorizable, code that can't be handled on the vector processor must be handled sequentially on the scalar machine. Naturally, the more code that can be vectorized, the greater the processing speed, but because 100 percent vectorization is unlikely in most applications, an extremely fast scalar machine complements the vector machine. A slow scalar machine can create a system bottleneck and degrade vector performance.

To get even more performance out of microprocessor technology, several vendors' designs use multiple processors that each incorporate vector and floating-point arithmetic units. Alliant was the first to use the technique at the minisuper level when it introduced the eight-processor FX/8 in 1985. The company upgraded the performance of the FX Series in February 1988 with the use of faster microprocessor technology. Convex went from a single-engine vector processor architecture to a multiprocessor architecture with the introduction of its C2 Series in March 1988. Convex' top-end system uses up to four parallel processors. Similarly, Gould uses up to eight processors in its NP1 Series. Parallel processors divide work among multiple processors to speed up job turnaround. But like vector processing, parallel processing has its limitations. Parallel processors work most effectively on independent data elements that can be split off to run on their own separate processors. However, jobs containing data dependencies—portions of code that cannot be run until earlier sections are completed—reduce parallel performance gains.

Proponents of the two newest architectures to come on the scene, VLIW and dataflow, claim their products overcome traditional limitations of vector and parallel processing. Multiflow and Cydrome contend vector and parallel machines will rarely attain peak performance because most applications cannot be completely vectorized or reformatted to run on multiple machines. That means there will always be portions of code that must be processed sequentially at much slower scalar rates. When a large portion of code must run sequentially, performance gains are minimal.

Both companies have developed machines that achieve parallelism through intelligent compiler technology. Multiflow, for instance, claims to run much more of a program in parallel by packing individual tasks such as adds, loads, and stores into very long instruction words (VLIWs) up to 1024 bits long. The compiler decides how best to restructure the code and pack each word without programmer intervention. These VLIWs are executed sequentially on a single CPU. Cydrome also uses a very long word. In many ways, the process is similar to conventional scalar processing. However, instead of carrying out only one operation using the customary 32-bit instruction word, the VLIW architecture can initiate up to 28 operations.

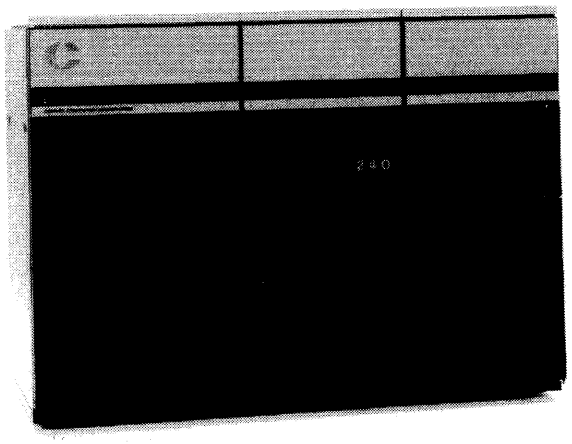

The Convex C240, the company's new top-end, four-processor system, uses scalar, vector, and parallel processing to attain a peak performance of up to 200 megaflops.

UNIX: The Tie that Binds

When it comes to architecture and hardware designs, minisupers come in many different flavors indeed. To make some of these somewhat unfamiliar approaches transparent to new technical users, virtually every minisuper vendor offers an implementation of AT&T's UNIX operating system. UNIX offers a common software platform and an environment familiar to technical users. The operating system has found believers at all hardware levels, from technical workstations, minisupers, and parallel processors to mainframes and supercomputers. Observers believe UNIX will begin to really take hold in both technical and commercial markets when a common industry standard is finally adopted. At the supercomputer level, most vendors offer proprietary operating systems and versions of UNIX. Cray led the way in 1986 when it announced UNICOS, an operating system based on AT&T UNIX System V. At the minisuper level, most vendors offer a UNIX implementation exclusively.

In an industry long dominated by proprietary operating systems and hardware incompatibility, UNIX presents users with some basic advantages. UNIX users can port applications based on UNIX fairly easily across UNIX processor environments of various sizes, from personal computers and technical workstations all the way up to mainframes and supercomputers. UNIX also provides program developers with a common platform on which to develop new applications across processor lines. Users don't have to become familiar with different command structures and syntaxes every time they move to a different system. Finally, the availability of third-party software makes the UNIX option particularly attractive.

For minisuper vendors selling to former VAX users, UNIX compatibility may not be enough. As already noted, many offer packages to help VAX users migrate applications to minisupers. These include software conversion tools that let program developers convert VMS/Fortran to a Fortran version that runs natively on a given

Minisupercomputers —Advancing Technologies

minisuper system. Virtually all vendors include VMS Fortran extensions and other popular extensions in their Fortran versions.

Before users can take full advantage of vector or parallel processing capabilities, applications and source code must be recompiled. Just about all vector systems, from Cray supercomputers down to mainframes and minisupers, offer Fortran compilers that automatically vectorize user source code. Alliant and Convex offer compilers that also search for parallel processing opportunities.

Multiflow and Cydrome, of course, don't use vector or multiprocessor approaches. Instead, Multiflow's compiler automatically analyzes code to determine the best way to pack instructions into very long words. Likewise, Cydrome's Fortran compiler examines code to implement dataflow technology. (For more information about specific minisuper technologies, please refer to the vendor product sections that follow.)

By necessity, all supercomputer vendors support Fortran, since it remains the language of technical computing and since most engineering/scientific applications are written in Fortran. Vendors also support C, Pascal, and Ada. C is becoming important because of the increasing use of UNIX, while Ada is becoming important because of Department of Defense requirements.

The presence of Cray and Digital Equipment in the market has encouraged the development of third-party software for scientific machines. Currently, more than 400 applications have been developed for Cray alone.

WHO IS SELLING MINISUPERS?

To better understand all the major products and vendors, Datapro has prepared detailed profiles of six minisuper companies, presented in alphabetical order. For a quick overview of each specific product offering plus price/performance comparisons, please refer to Tables 1 and 2.

Alliant Computer Systems Corporation

In the relatively short history of the minisuper computer, Alliant and Convex have emerged as the two industry leaders. Although the two start-ups took two distinctly differing paths to find their way to the top, both companies have nurtured sizable user bases within the last three years. Alliant has installed more than 200 systems since shipping its first system in 1985. The company was also in good financial shape as it entered 1988. The Littleton, Massachusetts-based firm reported a 75 percent increase in revenues in 1987 over 1986 and a 51 percent increase in net income. Revenue for 1987 was $53.8 million, compared to $30.8 million for 1986. Income for the year was $3.6 million, compared to $2.4 million for the year before.

Now that more competition has joined the minisuper fray, acquiring new customers and holding on to their

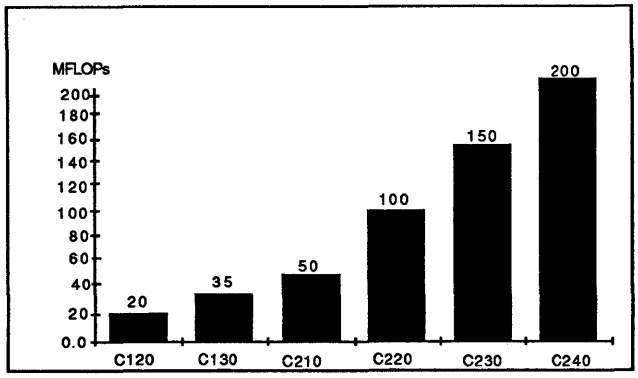

Convex moved from a single-processor architecture to a parallel architecture with the March 1988 announcement of the C Series. The chart graphically represents peak-performance increments among the six models in megaflops.

respective leads will be big challenges for Alliant and Convex. More competition has led to price-cutting and discounting.

Alliant responded to new competitive pressures, first with the October 1987 announcement of the FX/4, a new low-end system, and then again in February 1988 with the introduction of the FX/40 and FX/80, second-generation minisupers. The FX/40 and FX/80 replace the FX/8. The Alliant line now consists of four field-upgradable models ranging from the FX/1 with a peak performance of 11.8 megaflops to the FX/80 with a peak performance of 188.8 megaflops. Previously, top-end peak performance was half this.

Alliant introduced the FX/4, with up to four parallel vector processors, to make its systems more affordable to low-end VAX users and departmental technical workgroups. The system has a peak performance of 47.2 megaflops and a starting price under $100,000 for a one-processor system. FX/40 and FX/80 systems, meanwhile, offer up to 2 times the performance of Alliant's first-generation machines, first delivered in 1985, and up to 15 times the performance of a VAX 8700 in computationally intense applications, the company claims. While the newest systems offer more raw horsepower than the earlier models, the product line continues to hover around the $100,000-to-$500,000 price range. At this price, Alliant is clearly targeting the VAX user planning to upgrade to a more powerful system. Archrival Convex, on the other hand, has lately raised its sights to target the high-end minisuper user with larger power needs and perhaps more money to spend. (Refer to the Convex section below.)

Alliant recently acquired Raster Technologies Inc., a supplier of color graphics systems. Under the terms of the agreement, Raster will become a wholly owned subsidiary. Alliant apparently plans to bring high-performance 2-D and 3-D graphics to its minisuper systems. At press time, Alliant would not disclose how Raster Technology would be integrated into the Alliant product line. Similar to

Minisupercomputers —Advancing Technologies

other minisuper vendors, however, Alliant has been encouraging technical workstation users to use Alliant minisupers as compute- and I/O-intensive servers. Using Raster graphics technology, users will be able to see the results of compute-intensive applications displayed visually.

From the start, Alliant has achieved near-supercomputer performance using parallel processing techniques combined with vector processing. To make the process transparent to VAX users and other supermini users more accustomed to conventional single-processor technology, Alliant introduced FX/Fortran, the first minisuper compiler to automatically analyze existing Fortran code for *both* vector and parallel opportunities. Additionally, the compiler contains VAX/VMS extensions and other popular extensions. This allows users to transport and recompile existing applications from VAX machines and other systems with little or no code rewriting. The intelligent compiler technology has surely helped Alliant establish an early lead in the minisuper market. The company supports FX/Ada, Ada, C, FX/C, Pascal, and Alliant Assembler. Lisp and APL are available through third-party vendors.

Alliant's product line continues to operate under Concentrix, an enhanced version of Berkeley UNIX 4.2. The Alliant version of UNIX was enhanced to support multiple forms of parallel processing, 256 megabytes of physical memory address space, 2 gigabytes of virtual address space per process, and a high-performance file system that includes parallel access to I/O devices.

On the hardware side, the FX Series features two processor types arrayed around a bus-oriented parallel architecture. Computational Elements (CEs), which have since been upgraded to Advanced Computational Elements (ACEs) on second-generation machines, concentrate on compute-intensive portions of user applications. CE/ACEs are high-performance, pipelined processors with integrated vector and floating-point instruction sets. Multiple CEs or ACEs can work in parallel on a single application. The newer ACEs are 64-bit processors with a peak capacity of 23.6 megaflops per processor, twice the rating of the original CEs. ACEs use 20,000 gate CMOS very large-scale integration gate arrays, while CEs use 8,000 gate arrays. Users with first-generation FX/8 machines can replace CE boards with plug-compatible ACEs. The top-end FX/80 supports up to eight ACEs, the same as the earlier FX/8, and the FX/40 supports up to four.

Interactive Processors (IPs) execute interactive tasks and the operating system in parallel with one another and help to maintain system responsiveness. IPs access global memory through multiple IP cache memories. On FX/40 or FX/80 systems, each IP contains 4 megabytes of local memory. IP cache size is 64 kilobytes on FX/40 models and 128 kilobytes on FX/80 models. The FX/80 supports up to 12 IPs and the FX/40 supports up to 6.

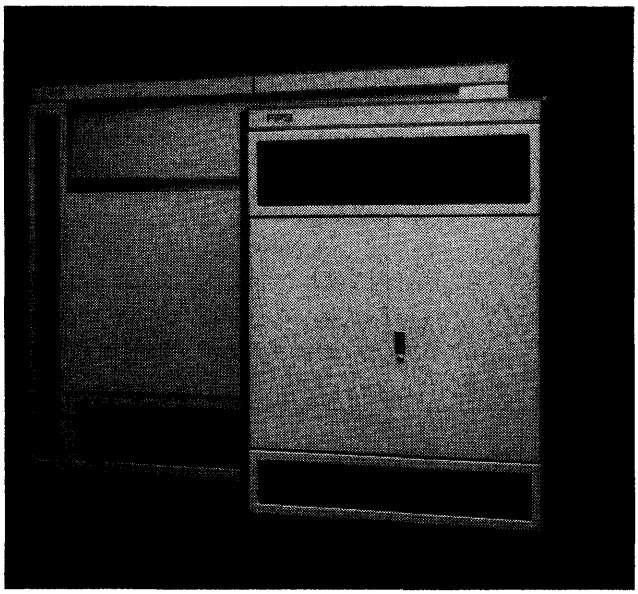

The Floating Point System M64 line of minisupercomputers attaches to Digital Equipment VAX systems, IBM mainframes, and other systems to improve the single-job turnaround time of numerically intense applications. Peak performance ranges from 6 megaflops at the low end to 341 megaflops at the high end.

Convex Computer Corporation

The Convex C1 minisuper, first introduced in 1985, was an industry success. Industry watchers wondered what the Richardson, Texas-based company was going to do for an encore. Convex responded on March 1, 1988 and scored another marketing coup. The company introduced a second generation of minisupers, which are 10 times more powerful than the original C1. The six-processor C Series models feature a performance range of 20 megaflops to 200 megaflops, overlapping the performance of a low-end supercomputer. The systems sell for between $275,000 and $1.35 million, still well within the price range of departmental engineering/scientific users, but below the entry-level price of a Cray supercomputer. The entry-level Cray X-MP/14se has a starting price of $2.5 million.

To achieve greater performance, Convex moved from single-processor technology to a parallel-processing architecture. These second-generation systems replace the original C1, making the Convex line more competitive with other recently introduced minisupers. As the apparent minisuper market leader, Convex must continue to make price/performance adjustments to maintain its lead within a market that has already become choked with new competitors. Convex' newest systems, featuring from one to four processors, are positioned towards the upper end of the minisuper market where there aren't as many competitors. The addition of five more Convex systems gives its user base a smooth upgrade path from a high-end Digital Equipment VAXcluster to an entry-level Cray supercomputer. Additionally, Convex no longer depends on one or two models. The new C models are said to be five times more powerful than a VAX 8700 and more powerful than

Minisupercomputers —Advancing Technologies

a Cray X-MP/14. At a well-attended New York City press conference in March, when the new models were announced, Convex President Bob Paluck suggested it may now be more appropriate to call Convex a *super* rather than *mini*super contender.

Favorable price/performance, VAX/VMS compatibility, and the availability of third-party software have been key Convex marketing strategies. Evidently, the strategies have been successful. During its relatively brief history, Convex has installed 260 machines at 150 customer sites worldwide. To build market share, the company has reached out to customers well beyond these shores. Internationally, the company has established subsidiaries in Canada, the United Kingdom, France, Germany, and The Netherlands and has established distributorships in Japan, Australia, Israel, Korea, Belgium, and Hong Kong.

For 1987, the company reported revenues of $69.6 million, an increase of 73 percent over 1986 revenues of $40.2 million. Net income for the year was $8.8 million, a 120 percent increase over 1986 net income of $4 million. International operations continued to lead revenue growth, the company reported.

New hardware includes the top-end C240, a four-processor, tightly coupled parallel "supercomputer" with a peak performance of 200 megaflops. The C230 features three CPUs and is rated at 150 megaflops. The C220 is a dual processor with a peak performance of 100 megaflops, and the C210 is a single-processor system delivering a peak performance of 50 megaflops. The C120 and C130 are both 40-megaflop entry-level systems. The C120 low-end system is a dressed-up version of the original C1 machine that launched the company three years ago.

With its newest machines, Convex reduced cycle time from 100 nanoseconds for the C1 to 40 nanoseconds for the C200 parallel machines and to 55 nanoseconds for the C130. The machines also use 10,000 gate array Emitter Coupled Logic (ECL) and 20,000 gate array CMOS semiconductor technology in scalar and vector units, while the earlier C1-based C120 machine uses 8,000 gate array CMOS and TTL semiconductor technology.

To achieve an efficient form of parallel processing, Convex introduced a job scheduling technology called Automatic Self-Allocating Processors (ASAP). In addition, the company has developed a new compiler that searches for parallel opportunities and automatically reformats user code to exploit these opportunities.

The ASAP concept tries to keep all processors within a parallel configuration busy most of the time. According to this scheme, a job can begin executing the scalar portions on a single processor, then branch to run on up to four processors during parallel portions of the job, and branch back to one or two processors as required. Convex contrasts this approach with static scheduling, which reserves an entire set of available processors for a single job. During scaler portions of the job, some processors will remain idle for some period of time, reducing system efficiency.

The new C Series machines continue to operate under Convex UNIX, based on Berkeley 4.2, an operating environment familiar to most technical users. Not surprisingly, many of Convex' customers are former VAX users who moved to Convex to attain better single-job turnaround on engineering/scientific jobs than was available on Digital VAXclusters. To make Convex systems even more attractive to VAX/VMS users, the company introduced COVUE, its Convex-to-VAX user environment, to ease migration to the Convex UNIX environment. COVUE retains familiar VMS environmental features, while providing users with greater performance. COVUE features COVUEshell (a VAX/VMS-compatible command language), COVUEedt (a compatible text editing facility), and COVUEnet (a DECnet-compatible networking facility). Additionally, COVUEbatch lets VAX customers use a Convex system as a computational server. Additionally, to make application portability from Digital and Cray environments transparent to users, the company features a Fortran compiler with VMS Fortran extensions.

Recognizing the need for applications software, Convex has worked to build a sizable software library. Currently, Convex offers more than 200 available third-party applications. Applications address such areas as geophysical research, computational chemistry, CAD/CAM/CAE, image processing, aerospace simulations, molecular biology, structural analysis, computational fluid dynamics, petroleum, electronics, and graphics. Available applications libraries include Veclib, BCSLIB, IMSL, LSQPACK, Math Advantage, NAG, Vectorpak, PCGPAK, and SMPAK.

Cydrome Inc.

The Cydra 5 Departmental Supercomputer, announced in January 1988, is one of the latest minisupers to enter an already crowded field. With so many start-ups selling minisupers, low-end supercomputers, parallel processors, and other high-performance systems to the same base of technical users, there doesn't seem to be room for any more start-ups. Even though it is a latecomer, Cydrome Inc. of Milpitas, California seems to have some good things going for it. The company's chief technical officer, Dr. B.R. Rau, designed an architecture using dataflow technology, an approach said to exploit a higher degree of parallelism than standard vector and multiprocessor approaches. Secondly, the company has formed a strategic partnership with Prime Computer Inc., an established vendor that will also market the machine as the MXCL 5. Cydrome, founded in 1984, can take advantage of Prime's worldwide distribution channels. Additionally, Prime now has a minisuper product that it can market to its CAD/CAM users who need to migrate to a system more powerful than Prime's current supermini offerings. Prime plans to market Cydrome machines to its U.S. and international manufacturing and financial users. Cydrome will directly sell and service the Cydra 5 in other U.S. markets outside

Minisupercomputers —Advancing Technologies

the Prime customer base. The company plans to concentrate on research, geophysical, and government markets. Cydrome has also completed a joint marketing agreement with Silicon Graphics, a producer of graphics workstations.

Cydrome is calling the Cydra 5 the first real commercial implementation of dataflow architecture, which the company calls Directed Dataflow. Similar to Multiflow, Cydrome says its approach is superior to conventional vector and parallel architectures because it achieves fine-grained parallelism. Fine-grained parallelism involves the execution of small overlapping operations such as adds, loads, and stores. Many competing products using multiple processors generally achieve course- or medium-grained parallelism, which divides a task up into larger coding chunks. Machines that rely on vector processing exclusively can only speed up portions of code that can be vectorized. The degree of vectorization varies, depending on the nature of the code and application.

Cydrome designers contend their Directed Dataflow architecture can optimize a broader variety of code, not just parts of code that can be vectorized. The machine can execute recurrences, unstructured parallelism, conditional branches, *and* vectors in parallel. Recurrences consist of repetitive calculations that need previous results to do current calculations. Unstructured parallelism refers to unrelated operations on unrelated data. Conditional branches typically occur in "if, then" instructions. Cydrome uses an intelligent Fortran compiler technology to achieve fine-grained parallelism. Many of the decisions directing the order in which the code will run are made at compile time, rather than during actual execution of the code.

At the hardware level, Cydra 5 consists of five subsystems: the Numeric Processor, Interactive Processors, Input/Output Processors, Memory Subsystem, and Service Processor. The Numeric Processor performs numerically intensive computing based on Directed Dataflow technology. The component contains multiple pipelined functional units, a context register matrix, and a unit for handling program branching. Interactive Processors handle the operating system execution and interactive work loads concurrently with the Numeric Processor. Input/Output Processors move data from peripherals and communications links. The Service Processor provides operator control and local and remote service access.

Applications run under the Cydrix 5.3 operating system, based on UNIX System V, Version 3, and the Cydrix Fortran 77 compiler, which supports VAX/VMS, IBM, and UNIX extensions. Under Cydrix and the Cydrix Fortran compiler, users can move existing applications over to Cydra with minimal code restructuring.

The Fortran compiler examines Fortran source code and reorders the coded instructions to maximize throughput. The compiler and Numeric Processor work closely to schedule program execution. Similar to Multiflow, Cydrome uses an instruction word up to 256 bits long, containing multiple operations. The processor and its multiple functional units can execute up to 56 parallel operations at a time. Up to seven new operations can be initiated each machine cycle. Processor cycle time is 40 nanoseconds.

The company plans to market the machine to departmental and technical workgroup users. Specific marketing areas include structural analysis, fluid dynamics, semiconductor design, mechanical design, computer simulation, earth resource management, computational chemistry, economics, and financial modeling.

Industry-standard applications available for the Cydra include Math Advantage, PLOT 10, Linpack, Itplack, FIDAP, Spice, a circuit analysis application, and Amber and Cedar, molecular modeling applications. Packages currently being ported include the NAG library; NAG tools; and Abaque, ADINA, and NISA, all finite element applications.

Current Cydra users include the Jet Propulsion Laboratory in Pasadena, California and Whitney-Demos Productions in Los Angeles.

Floating Point Systems Inc.

Floating Point Systems (FPS) has been selling minisupercomputers since 1981, well before the emergence of Convex and Alliant, and well before anyone was calling high-performance, 64-bit systems *minisupers.* "We were here first. We now have more machines in the field than anybody," a company spokesperson said. That's been the message Floating Point has been trying to hammer home in a marketing campaign designed to win back market share.

Floating Point, a Portland, Oregon company founded in 1970, first came to industry prominence selling board-level devices and then full-blown systems that attach to existing host systems. The company introduced its first 64-bit machines, the 164 Series, in 1981. In 1986, the company introduced the M64 minisupercomputer line, based on the original 164. Out of 8,500 FPS installations, the company classifies about 400 machines as minisupers.

The M64 line includes Models 20 through 60, machines which use scalar technology, and the Models 140 and 145, which use matrix technology. Performance ranges from 6 megaflops for the low-end M64/20 to 341 megaflops for the top-end M64/145. The product line attaches to Digital Equipment and IBM hardware and Sun and Apollo workstations. In addition to the minisuper line, FPS continues to offer the MP32 Series and 5000 Series of attached array processors. A third product line, the T Series massively parallel supercomputer, is now in limbo. The company has announced it will stop further development of the system unless it can find a partner to help fund the effort. FPS has installed 10 T Series systems since 1985 and will continue to support them.

Minisupercomputers —Advancing Technologies

The M64 minisuper line operates under the System Job Executive (SJE) operating system and includes optimizing Fortran (APFTN64), C (APC64), and Assembler (APAL64) compilers. Users entering jobs on VAX systems or other machines can run applications without change. The Fortran compiler contains VAX/VMS extensions. APFTN64 and APC64 execute on the front-end host or optionally on most FPS minisuper models.

Except for a few competitors, Floating Point had a large part of the market to itself until about the mid-1980s. That's when start-ups, such as Convex and Alliant, began selling to traditional Floating Point customers, typically VAX users who need more power to handle numeric-intensive applications. Today, more than two dozen vendors are competing for these same users. During the last few years, Floating Point watched its market share shrink and its profits plummet. In 1985, the company reported $126 million in sales revenue. In 1987, the company reported $84.3 million in sales and reported an earnings loss of $29.2 million. The company blames increased competition and pricing pressures for the loss. "We took some pretty heavy lumps," a spokesperson said.

To turn its fortunes around, the company brought in a new management team, closed an Ireland plant, and reduced the work force from a high of 1,600 employees in 1985 to the current level of 800. To win back market share, the company plans to refocus on its existing markets and reemphasize product advantages over its newest competitors.

Part of this effort involves convincing customers that it legitimately competes in the minisuper market. Floating Point is best known for manufacturing specialized processors that attach to a Digital VAX and IBM and Unisys 1100 Series mainframes, while Convex, Alliant, and others market so-called standalone systems. These days, Floating-Point prefers to call its machines "integrated" minisupers. While Floating Point concedes the standalone approach presents some valid advantages, the company contends the integrated approach can be more beneficial, particularly to VAX users. By attaching a Floating Point system to a VAX configuration, VAX users can continue to work in the familiar VAX environment and run their applications unchanged. VAX users purchasing a standalone system, on the other hand, will have to migrate VAX applications to a new environment. Convex, Alliant, and others offer VAX conversion tools and VAX/VMS emulation to ease the migration, but Floating Point contends VAX users can avoid the cost and disruption of installing a new machine simply by installing a high-performance attachment.

In many instances, VAX customers are happy with the performance of their installed VAXs doing commercial and interactive work, FPS contends. They only need to improve the single-job turnaround time of compute-bound production jobs. They may also find that large production jobs degrade interactive performance and make overall system performance sluggish. To eliminate these problems, users install a FPS model as a VAX back-end machine. In a typical FPS/VAX configuration, the front-end VAX handles interactive loads and general-purpose applications, while the back-end FPS can devote its full resources to handle entire numeric-intensive jobs. Using this approach, users can realize a performance improvement by a factor of 10 or better, FPS says.

Marketing FPS machines to VAX users is good for Floating Point and also good for Digital Equipment, since the attachments discourage users from replacing their installed VAXs. Recognizing that the two product lines complement each other, Digital and FPS entered into a joint marketing agreement more than two years ago. Under the agreement, Digital salespeople can offer customers FPS/VAX packages. Before the agreement, users had to order the machines separately from each company.

Besides Digital Equipment, FPS has agreements with third-party software vendors. Popular applications that run on M64 machines include Abaqus, Ansys, Marc, and MSC/Nastran for structural analysis; FL057 and ARC3D for transonic aerodynamics; FIDAP for finite element analysis; CPS1, Eclipse, Pores, Vespa, and VIP for petroleum and geophysical applications; Qspice and other Spice derivatives for circuit design; and Discover, Amber, and Gaussian 86 for computational chemistry. Additionally, FPS has entered a new area, financial modeling, with the delivery of an M64/60 minisuper to the Wharton School of Business of the University of Pennsylvania.

Multiflow Computer Inc.

Of all the recent minisuper start-ups, Multiflow Computer has attracted a lot of attention, and for good reason. Within the last year, the Branford, Connecticut company has introduced a minisuper line that speeds up engineering/scientific applications without relying on the usual multiple processor and vector processor approaches. Users don't even have to rewrite existing code to tailor it to the new machine. In fact, the company claims its TRACE line has been capable of running existing "dusty deck" code transported without modification from other systems. This lets users realize big performance gains without the expense of costly reprogramming and programmer retraining.

The four TRACE systems announced to date have a performance range of from 41 MIPS at the low end to 215 MIPS at the top end. They can also perform 11 to 60 million floating-point operations per second in full precision. The three available systems, the TRACE 7/100, 7/200, and 14/200, sell for between $197,500 and $399,500. No delivery date has been announced for the 28/200 top-end system. In the brief time the 7/200 model has been on the market, it has outperformed Convex C1 models and performed competitively against 3090 mainframes configured with Vector Facilities. Performance ratings in this case were derived from standard Linpack benchmarks. At the time the benchmarks were done earlier this year, Convex' newest systems had not been announced.

Minisupercomputers —Advancing Technologies

With features and performance numbers like these, it's easy to understand why the company has found so many early believers. In less than a year, the company has installed 33 systems (as of the first quarter of this year). Customers include Sikorsky Aircraft, Grumman Data Systems, Motorola Inc., Semiconductor Products Sector, Procter & Gamble, and Yale University.

Multiflow achieves near-supercomputer performance using a very long instruction word (VLIW) architecture combined with intelligent compiler technology. The company's three founders, Dr. Joseph A. Fisher, John J. O' Donnell, and John C. Ruttenberg, first developed these technologies while at Yale University and the Courant Institute of Mathematics. Unlike competing minisupers, the approach does not require multiple processors or specialized vector hardware. VLIW TRACE systems actually execute instructions sequentially using only a single CPU, much like a conventional single-processor supermini.

The real speedup comes through executing multiple instructions at one time packed into one long word. Each long word may contain many different logical units such as add, multiply, divide, load, and store instructions. By handling multiple operations in one program block, the machine achieves a kind of internal parallelism through software. The two entry-level TRACE systems pack seven different operations into each 256-bit word, while the TRACE 14/200 packs up to 14 instructions in a 512-bit word. The top-end TRACE 28/200 packs 28 operations into a word 1024 bits long. Compared to competing systems that execute one operation using standard 32- and 64-bit words, 1024 bits is a very long word indeed.

To make the VLIW approach work, Multiflow developed the TRACE Scheduling compacting compilers, a compiler technology that examines user code to determine the best way to pack and overlap operations before forming big instruction words. While simple load and add instructions can be handled fairly easily, VLIW machines traditionally have had problems handling conditional branch instructions such as "if, then" statements. Branch instructions limit the number of operations that can be handled as a block for simultaneous execution. Since conditional branches are common to scientific code, this can be a big obstacle. According to Multiflow estimates, conditional branches occur every 7 to 10 lines in typical scientific code.

Multiflow overcomes this problem by first predicting what branching operations are likely to occur based on statistical probabilities. Predicting the outcome of a conditional branch is not as chancy as it sounds, since conditional code tends to branch the same way most of the time. When the compiler makes its predictions, conditional branches are packed with other instructions in a long word. The compiler then processes the remaining code using VLIWs until all the code is processed. Should the compiler make an incorrect guess about a conditional branch, the system puts in "compensation code" to correct the mistake. While VLIW approaches have been attempted before, the company is calling its so-called compensation code true "breakthrough" technology, since this proprietary feature overcomes previous VLIW drawbacks and makes the approach commercially practical.

The Trace Scheduling compiler includes extensions for VAX/VMS Fortran, Fortran 8X, and standard features from IBM Fortran compilers. The C compiler implements the description of C given in the C Programming Language. The compiler technology automatically takes advantage of the long word Multiflow architecture. This means users can migrate applications from other systems without having to rewrite code. Multiflow systems run under TRACE/UNIX, an enhanced version of Berkeley 4.3 UNIX with extensions and performance improvements for scientific and engineering users.

Multiflow, a privately held company founded in 1984, introduced its first product, the TRACE 7/200, in April 1987 and made its first shipments in July 1987. The company shipped five 7/200s during the third quarter of last year and eleven more during the fourth quarter. Systems installed during the first quarter of this year bring the total to 33. In February of this year, the company continued its momentum with the introduction of the 7/100 and 14/200. The 14/200 is said to have two times the peak of 7/200. The new entry-level 7/100 is said to have two thirds the peak performance of the 7/200.

Customers who have installed TRACE machines are using them for integrated circuit design, finite element analysis, computer-aided engineering, fluid dynamics, wave propagation, numerical analysis, image and signal processing, and other applications. Available third-party applications include Abaqus, Ansys, FIDAP, Gaussian 86, IMSL, Math Advantage, NAG, Nastran, Spice, and others.

To build market share, Multiflow has formed alliances with Apollo Computer and Silicon Graphics, two scientific workstation vendors, and has also established overseas distributorships with GEI Rechnersysteme of West Germany and Ing. C. Olivetti & C., S.p.A. of Italy. GEI will distribute systems in West Germany and nonexclusively in Switzerland. Olivetti distributes systems throughout Western Europe, excluding West Germany and France.

Scientific Computer Systems Corporation

While most minisuper vendors have set their sights on the VAX/VMS user and similar technical users, Scientific Computer Systems (SCS) has focused attention on users who plan to someday migrate to a Cray supercomputer. Many such users have purchased time on a Cray through a time-sharing service bureau and have developed Cray-compatible applications. SCS machines execute Cray X-MP/24 or X-MP/416 instruction sets and can handle Cray applications with little or no reworking. Typically, SCS targets users who would like to purchase a Cray, but can

Minisupercomputers —Advancing Technologies

only afford a departmental supermini. They're part of what SCS calls the "Cray proximity market."

Although a Cray strategy appears to have much potential, SCS minisupers have not sold as well as similar systems from Convex and Alliant. SCS shipped its first machine in July 1986 and has since sold about 30 systems worldwide. By contrast, Convex has sold more than eight times that number since shipping its first machines in 1985. Expanding market share will continue to be an uphill battle, since more competitors have come on the scene. Up to now, SCS has been one of the few vendors to seek customers at the high end of the minisuper price/performance range, where there haven't been as many competitors. SCS has since been joined by Convex, which has taken performance to the entry-level supercomputer level. "We were there first," an SCS marketing spokesperson noted. "Convex has caught up to us."

To strengthen its marketing position and reverse an otherwise slow start, SCS brought in a new sales and marketing team last year, announced a new entry-level system to differentiate its product line, upgraded existing hardware, and introduced SCENIX, an implementation of UNIX System V Release 3.

Heading the new management team is Barry Rosenbaum, who became the company's new president and chief executive officer in 1987. Rosenbaum was formerly vice president of international sales at Convex. Steve Campbell, also formerly with Convex, joined SCS earlier this year as vice president of marketing.

In January, SCS introduced the entry-level SCS-30/XM, a system well within the price range of a VAX/VMS user. The system delivers 75 percent of the performance of an SCS-40, the company's first product, and sells at 60 percent of the price. The SCS-30 has a peak performance of 33 megaflops and sells at a base price of $371,000. It became available during the second quarter of this year. The SCS-40, first introduced in July 1986, has since been upgraded to the SCS-40/XM. The model has a peak performance of 44 megaflops and has a base price of $595,000. The SCS-40 performs at 25 percent capacity of the entry-level Cray X-MP/14, which sells for $5.5 million.

Both systems continue to be the only minisuper products that feature full compatibility with the popular Cray X-MP Series. In addition to executing Cray instruction sets, they continue to support Cray CTSS and COS operating systems. The addition of SCENIX, an operating system based on UNIX, now addresses what had been a glaring product weakness. SCENIX contains key extensions to make it compatible with UNICOS, Cray's implementation of UNIX. SCENIX became available during the first quarter of this year.

To enhance its appeal among VAX/VMS users, SCS also offers SCS/VMS compatibility tools that help VAX users migrate to SCS systems. Similar to products offered by competing minisuper vendors, the SCS compatibility tools let former VMS users move their applications over to the new system while also letting them continue to work within a simulated VMS environment. CTSS/VMS Command Language software contains modified or newly created VMS commands that let users perform CTSS tasks. An SCS library of data conversion routines converts VMS data files to SCS binary data files. A third product, the CODETRAN translator, automatically translates existing VMS Fortran source code to CFT Fortran source code. With minimal hand translation, the converted code will compile and execute on SCS systems.

At the moment, SCS offers more than 150 applications, many of which are Cray compatible. The second edition of the SCS Applications Software Catalog, released during the first quarter, features programs for structural analysis, computational chemistry, computational fluid dynamics, nuclear engineering, general-purpose graphics, mathematics, and statistics. Additional application software, particularly in the areas of petroleum and seismic processing, will be available in the third edition of the catalog, scheduled for release in July.

To remain competitive in the software area, SCS continues a major marketing and development alliance with Boeing. Boeing has set up a benchmarking center for SCS systems in the Seattle area and is supplying applications libraries. Boeing also helped SCS develop SCENIX.

At the hardware level, SCS now offers the Extended Memory (XM) feature. Introduced in December 1987, the Extended Memory option lets users add a gigabyte of memory. Included in the option is an upgrade from the Cray X-MP/24 instruction set to the functionality of the Cray X-MP/416. The option also implements gather/scatter and compress index instructions on the SCS-40/XM. These are advanced vector instructions that provide support for high-speed indirect addressing of array elements. The larger memory also reduces I/O bottlenecks, since more data can reside in real memory. Job swapping will operate at faster memory speed, which is 14 times the speed of disk I/O, according to SCS. The feature will benefit memory-intensive applications including computational chemistry, structural analysis, computational fluid dynamics, and other applications that use large data structures in their computational models, according to SCS.

An SCS spokesperson said the company plans to announce new products in future months to broaden the company's market appeal. At some point, for instance, the company will introduce new hardware to let users load programs onto SCS machines directly. Currently, users input jobs through a front-end VAX. □