SPECIAL ISSUE

SUPERCOMPUTERS: THE PROLIFERATION BEGINS

Now the hottest computer game around, growing at nearly 40% annually, they include a rapidly widening collection of products, large and small; with the current flood of introductions, they range from the traditional monsters to single-user 'desktop' models

by Tom Manuel

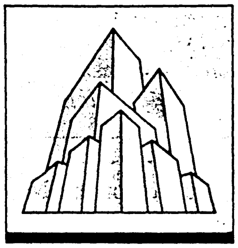

Supercomputers are proliferating all up and down the price and performance spectrum, in essence redefining the nature of supercomputing. The field now spans a range from single-user systems that may deliver nearly the full power of the first modern supercomputer, the Cray-1, to machines with almost 25 times the Cray-1's power. New circuit technology, as well as advances in hardware and software architecture, is driving prices down, making it easier for those who want a supercomputer to buy one. For under $1 million, users can buy a machine with the same performance of the previous generation's largest supercomputer. So the installed base is exploding as potential users realize that the number-crunching function of supercomputers can be applied beyond scientific and military applications: to product design, business, and much more.

And the number of supercomputer suppliers is exploding, as well. The pace of new product introductions has been furious in 1988's first quarter. Four are introduced in this special issue alone. At least five other introductions were unveiled in recent weeks, and several more are expected soon. That would mean adding more companies to the 31 vendors that currently have supercomputer products (see table, p. 53). And applications are taking off as software, much of it based on the burgeoning number of parallel-processing architectures, becomes available (see p. 75).

But as good as the news is, some analysts are cautioning that there are so many companies scrambling in the minisupercomputer and midrange market that a shakeout may occur. It could even be under way. The spectre of Japanese competition in the world market also hangs over the U. S. supercomputer industry. Although a trend to parallel architectures is clear in the U. S., the Japanese are sticking to single-processor technology (see p. 57). In the U. S., the Japanese have not jumped the hurdles necessary to make inroads and even on their home turf, Cray Research Inc. of Minneapolis still holds its commanding position.

Japan is not the only threat to the supercomputer market's upstarts, however. Everyone is waiting to see what the two largest U. S. computer companies, IBM Corp. and Digital Equipment Corp. (see p. 56), are going to do in supercomputers. How the marketing pie is eventually divided will depend on many factors, including how the spate of available technologies and systems eventually match up to application areas.

One way to look at the computer spectrum is as a triangle with three layers, says Steven Chen, founder and president of Supercomputer Systems Inc., Eau Claire, Wis. (see p. 54). Supercomputers are at the apex, mainframes in the middle, and work stations—the broadest market and most competitive—at the bottom. "Historically, as supercomputer vendors have extended

NEW TYPES SHOULD FUEL GROWTH

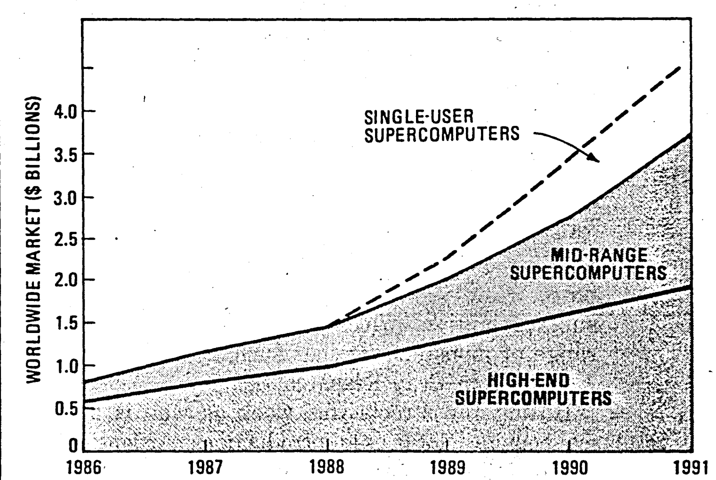

SOURCE: (AVERAGE OF DATA FROM) HAMBRECHT AND QUIST AND *ELECTRONICS*

SUPERCOMPUTERS—WHAT'S IN A NAME

Supercomputers are those computers that are the most powerful in existence at a given time—such was the old definition. Today, the proliferation of machines capable of high-speed floating-point operations up and down a broad performance range requires a revamping of that definition.

As the product spectrum broadens, the type of application software used may be the best way to define a supercomputer. Even if a computer fits on a desk and is dedicated to one user, it can still be called a supercomputer if it is designed for high-speed, calculation-intensive programs with an insatiable hunger for floating-point peformance, though it may offer only a fraction of a Cray-1's power.

Therefore, all machines specifically designed for high-speed floating-point calculations can be called supercomputers. Within this definition there are three basic types of architectures: single-processor machines; minimally parallel computers with a few processors; and a variety of highly parallel architectures. Examples of the first kind are the original Cray-1 and the current Japanese supercomputers. The second grouping includes the Cray X-MP and most other high-end systems, midrange supercomputers, and single-user supercomputers that use an architecture based on a small number of the fastest possible processors in a shared-memory multiprocessor configuration. The third supercomputer class contains a family of highly parallel multiprocessor systems configured as hypercubes or with several other interconnection topologies. Examples here would be the Connection Machine from Thinking Machines Inc. and the Megaframe Supercluster from West Germany's Parsytec GmbH.

Most supercomputers also have a built-in vector processor to complement the standard scalar processor. A vector processor is a computer which uses a pipeline processing technique to execute the same instruction on the entire vector, which is a list of numbers. Scalar computers execute instructions one at a time on individual values rather than vectors. The best supercomputers strive for a balance in performance since typical computationally-intensive applications programs do some of each. *—Jonah McLeod.*

the tip of the triangle further and further, technology developed at the top invariably moves down into the broader marketplaces below," Chen says. "What happens [in high-end supercomputers] tends to indicate the direction of the whole [computer] industry. When we move the tip, it creates a vacuum in between." It is this gap that's attracting minisuper makers, he adds.

The supercomputer market itself is segmenting into three levels—based primarily on price and performance. A new segment, single-user supercomputers, or the "supercomputer on a desk," is expected to match or exceed the growth of the other supercomputer categories (see figure, p. 51). The two existing segments, high-end supercomputers and midrange supercomputers, are also expected to grow considerably faster than the overall computer market—40%, as opposed to the 10% growth for computers on the whole.

Growth at the high end of the supercomputer market and in high-end work stations are both continuations of trends. But something else may be happening. "The one really profound change that's occurring is in between—in the minisupercomputer market," says Carl Ledbetter, president and chief executive officer at ETA Systems Inc., St. Paul, Minn.

The proliferation of machines has really hit an upswing this winter. It was capped this week by Convex Computer Corp.'s big announcement of the C series (see p. 59), a family of systems based on a second-generation processor. The Richardson, Texas, company ushered in the age of affordable supercomputing with its 1984 introduction of the C-1 minisupercomputer. Now it is taking the wraps off its first big supercomputers, which push into the performance range of monster machines such as the low-end X-MP models from Cray Research Inc.

Hitting the streets this week as well is a pair of eye-popping single-user introductions. Ardent Computer Corp. unveiled the graphics-based Titan (see p. 65) and Apollo Computer Inc., the well-known maker of engineering work stations, brought out a Domain Network-compatible machine called the Series 10000 (see p. 69). The fourth new supercomputer introduced in this special issue is a highly-parallel computer, which uses the powerful Inmos transputer, from West Germany's Parsytec GmbH (see p. 71). Faster machines need faster communications and Integrated Photonics Inc., Carlsbad, Calif., is trying to fill that need with a new very high-speed fiber-optic link for connecting supercomputers (see p. 74).

Other recent product introductions include several from war-horses in the supercomputing arena. These include the Y-MP from Cray Research Inc.; two machines (Trace 7/100 and Trace 14/200) from Multiflow Computer Inc.; the FX/40 and FX/80 systems from Alliant Computer Systems Corp.; Elxsi's model 6460; and the Cydra 5 and MXCL 5 from Cydrome Inc. and Prime Computer Inc., respectively.

Coming up soon are at least two other new machines. Stellar Computer will soon unveil its single-user graphics supercomputer and Scientific Computer Systems Corp. is readying for introduction later this year its second-generation machine (see p. 73). Analysts are also watching Multiflow, which jumped into the fray for midrange supercomputers late. But Multiflow's Trace 7/200 computer took off right away and the Branford, Conn., company is concentrating on building its sales and marketing forces and attracting important third-party application software.

One big question is how many of the newcomers will survive. Six companies are already missing from the action and either failed to get a product to market or were unable to stay in business: American Supercomputer, Celerity, CHoPP Computer, Culler Scientific, Denelcor, and Vitesse Semiconductor's computer division.

"Early 1988 is a critical period as many companies scramble to introduce products," says Jeffry Canin, a technical analyst at Hambrecht & Quist in San Francisco. He says venture capitalists fueled the midrange su-

THE PROLIFERATION: MORE THAN 60 DIFFERENT MACHINES NOW ON SALE

	Vendor	Product(s)	Peak Megaflops	Maximum No. processors	Linpack[1] Benchmark 100 by 100 (mflops/proc)	LPB[2] 300 by 300	No. of systems shipped or first ship date	Primary technology	System price range	No. of major application packages
High-end supercomputers	Cray Research	X-MP	1,200	4	39	480	123	ECL	$2.5 to $16M	500
		Cray-2	1,800	4	18	93 1p	11	ECL	$20M	150
		Y-MP	2,400–3,600	8	na	na	Q3 88	ECL	$12 to $17M	150
	ETA Systems	ETA 10-P	750	2	25	80 1p	1	CMOS	$995,000	16
		ETA 10-Q	947	2	31	101 1p	na	CMOS	from $1.2M	16
		ETA 10-E	6,857	8	56	182 1p	2	CMOS	from $5.5M	16
		ETA 10-G8	10,286 (e)	8	84 (e)	273 (e) 1p	Q4 88	CMOS	from $8.9M	16
	Fujitsu (Amdahl in U.S.)	VP-30E	220	1	na	na	16	ECL	30M Yen/mo. to 70M Yen/mo.	20
		VP-50E	280	1	14	100	16	ECL		20
		VP-100E	425	1	16	159	10	ECL		20
		VP-200	850	1	17	220	4	ECL		20
		VP-400	1,700	1	na	na	2	ECL		20
	Hitachi	S-810/5	160	1	na	na	20 total	ECL	40M Yen/mo. to 57M Yen/mo.	21
		S-810/10	315	1	na	na		ECL		21
		S-810/20	630	1	17	158		ECL		21
		S-820/60	1,500	1	na	na	1/88	ECL		21
		S-820/80	3,000	1	na	na	1/88	ECL/BiCMOS	80M Yen/mo.	21
	IBM	3090/VF 6 models	na	6	12	33 1p	200	ECL	$370,000 1st VF $230,000 ea. 2 thru 6	na
	NEC (HNSX in U.S.)	SX-JA	250	1	na	na		CML	28M Yen/mo. to 64M Yen/mo.	100
		SX-1EA	330	1	37 (e)	224	14 shipped	CML		100
		SX-1A	665	1	42 (e)	na	18 contracts	CML		100
		SX-2A	1,300	1	43	347	All models	CML		100
Mid-range supercomputers	Alliant	FX/8	94	8	1.6	27	156	CMOS/ECL	$100,000 to $500,000	110
		FX/4	na	4	1.6	12	4	CMOS/ECL		110
		FX/40	na	4	2.3	na	2/88	CMOS/ECL		110
		FX/80	189	8	2.3	na	2/88	CMOS/ECL		110
	Convex	C120	20	1	3.7	15	260	CMOS	$249,000	>200 for all models
		C130	36	1	7.3	23.8	3/88	CMOS/ECL	$359,000	
		C210	50	1	10	32.8	3/88	CMOS/ECL	$549,000	
		C220	100	2	10	na	Q3 88	CMOS/ECL	$820,000	
		C230	150	3	10	na	Q4 88	CMOS/ECL	$1.07M	
		C240	200	4	10	126 (e)	Q4 88	CMOS/ECL	$1.3M	
	Cydrome/ Prime	Cydra 5/ MXCL 5	50 50	1 1	10.4 10.4	na na	6	ECL	$550,000 to $950,000	8
	Elxsi	6410	40	12	0.6	na	over	ECL	$140,000 (proc. only)	65
		6420	120	12	1.5	na	100	ECL	$395,000 to $2.6M	65
		6460	na	10	10	na	Q4 88	ECL	$695,000 to $4M	65
	Floating Point Sys.	FPS M64/60	38	1	5.9	33	na	ECL	$475,000 to $1.2M	200
	Gould	NP1	320	8	1.7	na	25	ECL	$395,000 to $2.9M	na
	Multiflow	Trace 7/100	11	1	4.2	na	2/88	CMOS	$197,500	15
		Trace 7/200	15	1	6	na	19	CMOS	$299,500	15
		Trace 14/200	30	1	10	na	5/88	CMOS	$399,500	15
	Saxpy	Matrix 1	1,000	32	na	na	1	CMOS	$900,000 to $1.8M	none
	SCS	SCS-40	44	1	8	37	30	ECL	$595,000	150
		SCS-30/XM	33	1	na	na	Q2 88	ECL	$371,000	150
	Sky	Vortex	10	1	2.8	4.8	200	CMOS	$10,000 to $50,000	none
Highly parallel computers	Ametek	Series 2010	80	512	na	na	Q3 88	CMOS	$45,000 to $495,000	5
	BB&N	GP1000	125	256	na	na	100	CMOS	$75,000 to $2.5M	<10
	Flexible	Flex/32	80	20	1	na	16	CMOS	$110,000 to $600,000	na
	Intel	iPSC/2 VX	1,028	64	na	na	1/88	CMOS	$325,000 to $1.57M	3
	Meiko	Computing Surface	1,000 or more	1,000 and up	na	na	more than 100	CMOS Transputer	$15,000 and up ($5,000 increments)	none yet
	NCube	NCube 10	500	1,024	0.135	na	na	NMOS/CMOS	$100,000 to $1.8M	20–25
	Pacific Cyber/ Metrix	HyperFlo	na	100	na	na	Q1 88	CMOS 68020	from $20,300	none
	Parsytec	Model 64	96	64	na	na	3/88	Transputer	$230,000 to $320,000	none
		Model 256	384	256	na	na	3/88	CMOS	$850,000 to $950,000	none
	Thinking Machines	Connection Machine	2,500	65,536	na	na	15	CMOS	$1M to $7M	none
Single-user super-computers	Apollo	Series 10000	140	4	6	na	Q3 88	CMOS/ECL	$70,000 to $80,000	1,800
	Ardent	Titan	64	4	6	24	5/88	CMOS	$79,000 to $150,000	na
	Stellar	GS1000	40	na	na	na	na	na	na	na

1: Linpack 64-bit precision, 100 by 100, all Fortran (mflops/processor)
2: Linpack 64-bit precision, 300 by 300, matrix-vector operations, Fortran (total system mflops, except where noted)
(e) Estimated
Lease prices available in Japan only

SOURCES: Company-supplied information and Linpack Benchmark Report No. 23 by Jack Dongarra, Argonne National Laboratory percomputer market and may have overloaded that segment with vendors.

The shakeout started in 1987 and will continue into 1988, says Peter Appleton Jones, president of Elxsi in San Jose, Calif. Of almost 40 companies trying to address the high-end computing market, he says, about half are newcomers who will be vulnerable because of immature product lines.

The most likely survivors of any shakeout would be Convex and Alliant, says Gary Smaby, technical analyst at Piper, Jaffray & Hopwood, Inc. in Minneapolis. Smaby stops picking at those two companies, however. One of the problems facing the midrange supercomputer vendors is that "the sales cycles have not shortened as expected," Smaby says. "But the bigger disappointment for the early players has been the margins. Price discounting in either competitive situations or by customers anticipating next-generation machines has set up a tougher price environment than we [analysts] expected this early in the market's history."

The high-end market may be more secure, says James R. Berrett, chairman and chief executive officer at Honeywell-NEC Supercomputers, Inc. "The big five has remained the same for five years: Cray and CDC (now ETA Systems, Inc.) in the U.S., and Fujitsu, Hitachi, and NEC in Japan," he says. "This business isn't conducive to startups because it takes $150 million in R&D to put a supercomputer on the drawing board."

In the U. S., high-end suppliers have watched over their shoulders to see where the Japanese are heading. The three big Japanese computer companies, Fujitsu, Hitachi, and NEC, have built very fast supercomputers using impressive hardware but they have not been able to sell more than a handful in the U. S., largely because of the federal government's actions to prevent the purchase of non-U. S. supercomputers.

For example, executives at HNSX Supercomputers, a joint venture startup between NEC Corp. and Honeywell, charge that the U. S. Department of Commerce stymied its efforts to sell an SX-2 system to the Massachusetts Institute of Technology last summer. "It is clearly acceptable for politicial barriers to stay in place while the technologies are comparable. But they will come down if American computers fall behind," Elxsi's Appleton Jones says.

And staying ahead of the Japanese will be a challenge. "Architecture and software are the keys to staying ahead in these markets," says Craig Mundie, vice president for marketing at Alliant Computer Systems. "You can't chase the ultimate megaflop via semiconductor technology. If you believe that we're already hitting some [semiconductor performance] limits, you have to believe that the real gains will come from architecture, where we have a fairly considerable lead."

Smaby of Piper, Jaffray & Hopwood agrees that the Japanese are playing catch-up in some areas. "The real advantage that U. S. companies will enjoy is software. That's the one barrier that they'll have to construct to keep out the Japanese," he says, "I think the Japanese are going to have a very difficult time establishing a foothold [in the U. S.] at the high end. Outside the U. S., they will be very competitive."

HOW STEVE CHEN PLANS TO BEAT SEYMOUR CRAY AT HIS OWN GAME

Steve Chen's drive to build the world's fastest supercomputer is putting his company on a collision course with Cray Research Inc., his former employer. Thanks to his agreement with IBM Corp., [*Electronics*, Jan. 7, 1988, p. 31], Chen apparently has his pick of Big Blue's technologies. That may come in handy because the task of Chen's company, Supercomputer Systems Inc. is to build a machine by the early 1990s that rivals Seymour Cray's proposed Cray-4 (see p. 55).

Chen is confident that starting his own company and teaming up with IBM was the right move. If he succeeds, the results will be impressive: a machine performing 128 billion floating-point operations/s—the same as the Cray-4. The system would use that power to support interactive, three-dimensional, real-time scientific modeling.

"This system will be so fast, there's no sense in talking about how fast it will do A plus B plus C," Chen says. "What you really want to know is how fast it will get a solution for an aircraft designer, say, who is sitting in front of this engine."

One competitor, Carl Ledbetter, president of ETA Systems Inc. and a former IBM development manager, is doubtful of the end results. "Chen's linkup with IBM is getting a lot more attention than the reality warrants," he says. "One of the cardinal rules about IBM is that if they talk about it, they're not serious." But if IBM stops talking, watch out. "If it turns into a real product effort, they'll go absolutely dead, stone-cold silent," Ledbetter predicts.

IBM sounds serious. "Chen's performance objectives are very aggressive—that's a good thing to back," says Irving Wladawsky-Berger, vice president for engineering and scientific computing in IBM's Data Systems Division. "We will go gung ho trying to help SSI build the best thing they can." The agreement's details are still up in the air, but essentially Chen gets access to IBM technology while the big giant gets marketing rights to the new machine.

Chen and his team of some 45 ex-Cray associates envision a system that will exploit parallelism to the hilt and use the most advanced technologies that the research laboratories of IBM and other potential partners have to offer. Code-named the SS-1, it could have anywhere from 32 to 256 processors, Chen says. It will probably be based largely on VLSI silicon bipolar circuitry, at more than 10,000-gate densities. Small- to medium-scale gallium arsenide devices may be used for high-speed drive and interconnect circuitry. Chen doesn't rule out the use of high-temperature superconducting technology based on new materials discovered at IBM. The SS-1 design team is also looking at fiber-optic lines for short-haul, high-bandwidth communication among the SS-1's processors and subsystems.

The system will also rely on new kinds of highly parallel, ultrafast imaging devices to emphasize data interpretation. By allowing scientists to spot problems, stop a program or simulation in midstream, and make modifications, Chen hopes the machine will allow scientists interact with 3-d simulation models while they are in progress. This faster problem-solving "will be almost like watching your VCR," Chen says. *–Wesley R. Iversen.*

e supercompu r company with the most weapons to fend off Japanese companies is Cray, with more than 60% of the world market. To stay ahead, Cray is keeping up its frenzy of new product developments (see panel, below). Cray's approach, however, is that it takes more than just hardware to be successful. "When you think about what's required to be successful, clearly, having the fastest machine is a tremendous advantage," says Marcelo A. Gumucio, executive vice president for marketing at Cray. "But having the best software in the world is also very important—the operating systems, languages, applications packages, networking capabilities, and so on."

Who does Cray consider serious competition overall? First NEC, because of its semiconductor capability, communications experience, and "deep pockets," Gumucio says. Number two on his list is "clearly IBM." Number three is ETA, which he says must improve its financial viability and develop sufficient software for its machine.

IBM is one of the two dark horses in the supercomputer race; DEC is the other. Supercomputer industry experts have been expecting DEC to introduce a vector-processing VAX machine to compete in the midrange supercomputer marketplace. IBM, on the other hand, has been offering a vector processing add-on to its top-end mainframe computers, the 3090 line, for two years now with considerable success.

What will be IBM's strategy for supercomputers between now and the early 1990's when it can sell the machine that Steve Chen's Supercomputer Systems Inc. is starting to develop? "The 3090 is our family into the 1990s, and then at some point in the early 1990s you can expect to see another family—probably 1990 or 1991," says Irving Wladawsky-Berger, vice president for engineering and scientific computing. IBM is also looking at parallel architectures beyond the six-processor 3090/600. Research at Cornell University and at IBM's Kingston, N. Y., facility aims to combine multiple 3090s. Vector processing could also play a big role in other parts of IBM's product line, says Wladawsky-Berger.

But ETA's Ledbetter views IBM's project with Chen's company as not being much different than any number of IBM projects with outsiders that are funded in the $10 million to $20 million range. Some develop into real products. Others don't. "That's IBM's way of stirring the pot and trying to see what develops," he says.

Most U. S. supercomputer companies have turned to a low level of parallel processing to increase performance. In most of these systems, the maximum number of processors is four. However, the trend is upward. For example, there are eight central processing units in the new Cray Y-MP and in the two biggest ETA machines, the top two Alliant models, the Gould NP-1, and the forthcoming Cray-3 has 16 processors.

However, these numbers are small potatoes compared to the number of processors typical in another class of machines, the highly-parallel computers.

THE TECHNOLOGY PUSHING CRAY'S THREE UPCOMING MACHINES

Cray Research Inc. isn't resting on its laurels after introducing the 3.6-billion-floating-point-operations/s Y-MP in February. Three new products are under development at its Chippewa Falls, Wis., laboratories.

There are two distinct product lines in Cray's stable. One is the X-MP and Y-MP, which continue the lineage emanating from the original Cray-1. The other is the Cray-3, an evolution from the Cray-2. With its very large directly-addressable physical memory, the Cray-2 was incompatible with the Cray-1. The next product in that family, the Cray-3, adds another differentiator between the two product lines, gallium arsenide circuits. The MP family will continue to use silicon technology.

The 16-processor Cray-3, scheduled for delivery as early as next year, will be the first major computer system built with GaAs logic. Ultradense packaging will reduce connection lengths and take full advantage of the machine's 200-to-300-gate GaAs logic chips. Featuring a 2-ns clock rate, it will be up to 10 times faster and less than half the size of its predecessor, the silicon-based Cray-2.

The Cray-3 will rely on low-voltage depletion-mode GaAs standard-cell circuits from Gigabit Logic Inc., Newbury Park, Calif. The system's 48,000 GaAs logic chips and 40,000 silicon memory chips are housed in 200 assemblies stacked in eight columns, each 25 assemblies deep. Every assembly houses four 16-by-16-in. power and logic plates, upon which 4-by-4-in. circuit boards containing the GaAs chips are mounted, with up to 1,024 chips per assembly (see figure).

SQUEEZING IN 88,000 CHIPS

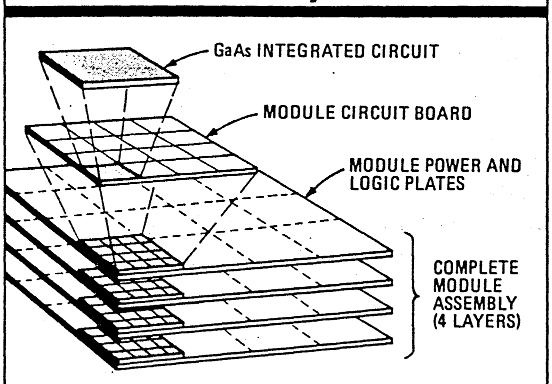

Gallium arsenide logic in the Cray-3 will use ultradense packaging to take advantage of 200- to 300-gate GaAs chips.

Meanwhile, the same team that developed the Y-MP [*Electronics*, Feb. 18, 1988, p. 31] is already working on its successor, known informally as the Y-Prime. It could ship by 1990.

The Y-Prime will be based on VLSI silicon logic. It will feature more and faster processors and a larger memory than the Y-MP. Cray is evaluating next-generation emitter-coupled-logic gate-array technologies that could push densities toward 7,000 to 15,000 gates, compared with the 2,500-gate-equivalent ECL arrays used in the Y-MP.

Furthest off is Seymour Cray's next machine, the Cray-4, expected between 1992 and 1994. The Cray-4 could collide with a system from Steve Chen (see p. 54). To be based on LSI GaAs circuitry, the Cray-4 is conceived as a 64-processor system using a 1-ns clock to achieve peak rates of 128 gigaflops, some 100 times faster than any machine currently available.

–Wesley R. Iversen

DEC's SUPERCOMPUTER ENTRY IS COMING—EVENTUALLY

Everyone is waiting for Digital Equipment Corp. to strike out against the minisupercomputer companies that have been invading its turf. The weapon, a VAX machine with an integrated vector-processing extension, was expected early this year. But it's probably still a year or more away, and rumors about DEC's progress continue.

"What DEC will do, that's the real excitement," says competitor Barry Rosenbaum, president of Scientific Computer Systems. "The long-awaited vector-VAX, however, is not expected until mid-1989, by my best guess. There is little doubt that whenever it comes, even if the technology is a half-generation behind, DEC can take away market share."

One analyst agrees that when DEC moves, it will uncork a winner. Gary Smaby, with Minneapolis-based Piper, Jaffray, & Hopwood Inc., says DEC has several development programs under way aimed at exploiting "new architectures and technologies used by minisupercomputer vendors. Within 12 to 18 months, we would expect to see some results. Our bet is that a competitive offering from DEC would take over market share leadership in another 12 to 18 months following the introduction," he says.

At DEC, executives try to dispel all the rumors. "We've made it clear that we'll have this capability [built-in vector-processing]," says Robert M. Glorioso, vice president for high-performance systems. But he won't be pinned down on when. And what he reveals about performance is scanty. The extension will allow "fairly respectable computing performance commensurate with what you can expect from a respectable ECL vector processor," Glorioso says.

DEC has a relationship with Floating Point Systems Inc., under which it sells the FPS minisupercomputer systems as attachments to VAXes. Digital recently extended that agreement despite FPS's problems.

This extension leads Jeffry Canin, senior technical analyst at Hambrecht & Quist, Inc., San Francisco, to speculate that DEC may have encountered technical delays in its vector-processing project. "Most of the success of Alliant and Convex has come at the expense of DEC," Canin says, "so they [DEC] will respond, but I don't expect an announcement for 12 to 18 months." Canin guesses that a vector processor from DEC would offer performance similar to that of the first generation of minisupercomputers, such as the original Alliant FX/8, with 94 megaflops peak performance, and the Convex C-1, with 40 megaflops peak.

Not everyone thinks that Digital is looking only at integrated vector processing to attack the supercomputer market. Omri Serlin, president of ITOM International Co., Los Altos, Calif., says DEC may be reevaluating the vector approach's pros and cons against the architectures created by Multiflow Computer Inc. and Cydrome Inc. These companies' architectures use very long instruction words that are an alternative to vector processing. "That's as good as and sometimes better than vector processing," Serlin maintains. "If I were DEC, I'd say, 'Hold it, let's look at other approaches'." *– Larry Curran.*

Many machines in this category are still considered research machines, although some, such as those from Sequent Computer Systems of Beaverton, Ore., and Encore Computer Corp., Marlboro, Mass., are strong in the commercial sector. Most of the other makers of so-called massively-parallel machines have added vector-processing capability or at least beefed up the floating-point performance in order to go after part of the supercomputer market. On at least some problems, they can boast very high peak performance.

The most massively-parallel computer, with 65,536 processors, is the Connection Machine from Thinking Machines Inc. This unusual machine has been reported to attain up to 7 billion instructions/s doing some specialized tasks, such as those where the same tasks are to be performed upon large volumes of data. However, such blinding speeds are not available for other applications until complex parallel programming is developed.

Another highly-parallel computer is the Butterfly GP1000 built by BBN Advanced Computers Inc. of Cambridge, Mass. It features up to 256 32-bit microprocessor-based CPUs in a multiple-instruction, multiple-data parallel architecture. "These are different from single-instruction, multiple-data machines like the Connection Machine," says Gary Schmidt, vice president of marketing. Other highly-parallel computers using the hypercube architecture include the iPSC family from Intel Scientfic Computers and the Ncube series from Ncube; the series 2010 made by Ametek Computer Research Division; and the Parts family from Computer System Architects. Highly-parallel computer makers with other variations in architecture include Flexible Computer, Meiko, Pacific Cyber/Metrix, and Parsytec.

One of the pioneering companies in the midrange of supercomputing, Floating Point Systems Inc., in Beaverton, Ore., had been quite successful in offering attached array processors to computer vendors and embedded-systems manufacturers and then graduating to stand-alone minisupercomputers. But it ran into trouble with its Hypercube machine, the T series.

Analysts say that the company's customer base and the market in general were not ready to deal with an exotic and incompatible massively-parallel machine built from a microprocessor that had no software to run on it. The only programming tool for the Inmos transputer chip offered when the T series was introduced was the Inmos-invented Occam parallel programming language, which was considered obscure and difficult. All this was detrimental to sales of the system. The company has taken steps to overcome these problems. And a new version, the T series Mark 2, should appear sometime in the second half of the year.

While the 31 manufacturers now competing in the three sectors of the supercomputer market are quite likely to shrink drastically in number over the next decade, it seems likely that by the year 2000, these classes of machines could be as dominant as today's standard varieties. □

Additional reporting by Paul Angiolillo, Lawrence Curran, Wesley R. Iversen, J. Robert Lineback, Jonah McLeod, and Larry Waller.