

# Changing of the guard leaves Cray out

## Washington

AN escalating shift from traditional supercomputers to massively parallel machines has left the industry's founder, Seymour Cray, with no products, no customers, and chilly prospects for success.

Two years ago, the titans of the supercomputer industry were assuring researchers that although new 'massively parallel' machines might be able to run test programmes with blinding speed, real-world problems required real computers. That means linear or 'vector' processors — the type of machines that have traditionally dominated the industry founded by Cray, who created Cray Research, Inc., in 1972 and then, 17 years later, left to begin Cray Computer Corporation.

Now, those same companies are designing massively parallel machines as fast as they can shift engineers from the traditional models. And Seymour Cray is out in the Minneapolis cold.

In December, after Cray missed a series of deadlines, the Lawrence Livermore National Laboratory cancelled its order for the first Cray 3, the machine on which Cray and his new company had pinned their hopes. Livermore also broke off an agreement to help Cray develop software for the new machine. With no other buyers on the horizon, Cray Computer announced last month that it was halting development of the Cray 3 and would try to produce a smaller and cheaper version of the same machine.

Even at Cray Research Inc., the word is out that the old ways will have to go. After ignoring the upstart massively parallel machines while it perfected the traditional vector processing supercomputers (epitomized by last year's C90 machine), the company has now made a somewhat tardy leap into the parallel world.

Last month, at a supercomputing conference in Paris, Cray Research announced that it was developing a series of massively parallel computers based on a new chip — dubbed 'Alpha' — from Digital Research Corporation. The new Cray computers will combine hundreds of Alpha chips — each one as powerful as the Cray 1, the machine that launched the supercomputer revolution 20 years ago. Officials hope to have by 1997 a machine that operates consistently at speeds at or above a teraflop (one trillion floating point operations per second).

Six of the seven leading supercomputing companies — including IBM, Fujitsu, Hitachi and NEC — are now working on massively parallel machines. They are riding on the coat-tails of Thinking Machines, Inc., the start-up company that introduced the first commercial massively parallel computer in 1986 and surprised nearly everyone by actually making it

work. Four of the five National Science Foundation (NSF) supercomputer centres now have at least one massively parallel machine, and half of those are from Thinking Machines.

This shift is due not so much to the success of massively parallel machines — which still lack sufficient software to solve most problems — as to the inability of traditional supercomputer manufacturers to improve their machines as fast as they once could.

While the massively parallel hardware is relatively simple — consisting largely of



Cray Research's C90: end of the line?

existing processors strung together by the hundreds or thousands — the hardware of traditional machines has become so complicated and temperamental that engineers who once improved performance by orders of magnitude overnight are now happy to achieve gains of a few per cent.

It is not for lack of effort. Both Seymour Cray and his protege Steve Chen, who left Cray in 1988 to form Supercomputer Systems Inc., have invested tens of millions of dollars into speedy gallium arsenide technology. But the technology has turned out to be more difficult than expected. Cray, for example, lost \$55 million last year struggling with the fragile chips, and neither company is expected to have its machines ready soon.

"I think Seymour and Steve Chen are barking up the wrong tree," says William Wulf, a University of Virginia computer scientist and former assistant NSF director for computer science. "There is no question about the demise of the behemoths."

Cray, Chen and the other traditionalists are up against nothing less than the laws of physics. Because the vector machines run one instruction after another, the only way to speed them up is to execute each instruction faster. With the speed of light and electric signals being an absolute, designers try to bring the parts as close together as possible and run them as fast as they will respond. The heart of the Cray 3,

for example, is no bigger than a shoebox and runs at more than 200 MHz. But warp speed and cramped components mean temperatures approaching the melting point for some parts. At least one vector machine has to run in a constant bath of liquid nitrogen lest it melt down.

In contrast, speeding up parallel machines is theoretically as simple as adding more cheap processors. Thinking Machines has made a computer with 64,000 processors, and even larger machines are available to the buyer with enough money.

Leading the endless quest for speed are what researchers call 'Grand Challenge' problems. These are attempts, such as climate prediction and DNA sequencing, to model reality on computers. "Because they come from nature," says Steger, "Grand Challenge problems have a high degree of parallelism in them." In other words, a lot of things happen at the same time in the real world, and keeping track of them takes a lot of processors.

But vector machines are still preferable for many of the problems for which researchers use supercomputers. Even for problems in which two-thirds of the operations could be conducted independently and simultaneously, traditional vector machine are orders of magnitude faster than parallel machines. Only when programs reach 98 or 99 per cent parallelism are massively parallel machines faster.

Few real-world programs reach that level. Most of the problems now being run with such impressive speed on massively parallel machines are what pundits call "embarrassingly parallel" — relatively unique simulations such as atmospheric modelling where the computer is required to recalculate conditions at thousands of unrelated data points. Although such programs are obviously important, the average program marches through its calculations one step after another — something massively parallel machines generally do no better than desktop PCs.

Because rewriting computer code to run on parallel machines is still a black art, experts predict that the traditional machines will last another decade. "The old guard isn't dead," says Bruce Steger, manager of the massively parallel programme at Cray Research. "Vector machines may be dinosaurs, but the dinosaurs lived for an awfully long time."

But even as it defends its old designs, Cray Research is working on a \$12.7-million grant from the Defense Advanced Research Projects Agency to hasten their demise. One goal is software that enables vector programs to run on parallel machines — just the thing to bury the traditional machines once and for all.

**Christopher Anderson**