Supercompeting Over Supercomputers

Prospect of a contest with Japan for preeminence sparks serious discussion of the government role

Last summer, Fujitsu, a major Japanese electronics and computer manufacturer, announced that it would market a computer six to eight times more powerful than the fastest American machines now available. This appeared to herald the end of a monopoly in high-performance computing held by the United States since the beginning of the computer era.

The Japanese challenge in supercomputers is powered by two ambitious national programs (*Science*, 17 December 1982, p. 1189), the National Super Speed Computer Project to develop a machine a thousand times faster than current supercomputers and the Fifth Generation Computer Project, an effort to develop computers that will embody artificial intelligence (AI) functions. Concern in the United States about the economic and national security implications of Japanese plans has prompted initiatives by both U.S. government and industry, but critics here question whether the response is adequate and some are calling for a U.S. national program in which the federal government would take a more aggressive role.

The United States currently leads both in the production of supercomputers—the Cray-1 and Cyber 205 machines have dominated the field—and in the development of hardware and software for future supercomputers. A catalyst for American concern about Japanese intentions, however, was a 1981 international conference convened by the Japanese in Tokyo to discuss and refine their plans for the fifth generation project. Western computer experts concluded that the Japanese were making a declaration that they would no longer be content with taking Western technology and improving on it, but were determined to seize the lead in innovative research and development.

The American response to the Japanese challenge has so far adhered to the pluralistic, decentralized style characteristic of American science. Particularly since the Japanese launched their two national programs, however, attention to the subject of supercomputers has picked up sharply here.

One U.S. reaction was expressed recently in a government-sponsored report* by a panel of computer experts, most of them from national laboratories and universities. The panel, headed by Peter D. Lax of New York University, looked at problems of access to supercomputers facing researchers in this country and at prospects for development of future supercomputers. A conclusion that attracted considerable notice was that "This panel believes that under current conditions there is little likelihood that the U.S. will lead in the development and application of this new generation of machines."

The Lax panel called for a "national program," but avoided specifying details. The panel did not suggest that the United States emulate the Japanese, but rather urged that the U.S. government do such things as improve access to scientific computing, boost manpower training in scientific and engineering computing, and support "research and development basic to the design and implementation of new supercomputer systems. . . ."

The federal government has exerted strong influence on high-performance computing since the end of World War II when it backed development of the ENIAC computer to compute ballistic trajectories. The federal role has been to serve as a principal customer for state-of-the-art, high-performance computers and as a patron for relevant basic research, mainly in national laboratories and universities.

That pattern continues. But within government, a debate is currently in progress over how the United States should respond to the challenge and what the federal role should be. Main parties to the discussion are the agencies most deeply involved with big computers. These are the Defense Advanced Research Projects Agency (DARPA), which manages the Defense Department's basic and applied research programs, the National Science Foundation (NSF), the National Aeronautics and Space Administration (NASA), and the Department of Energy (DOE).

DARPA has been a consistent and generous supporter of research in microelectronics and computers. It took the lead in funding AI research and contributed significantly to development of the three principal university centers in AI—

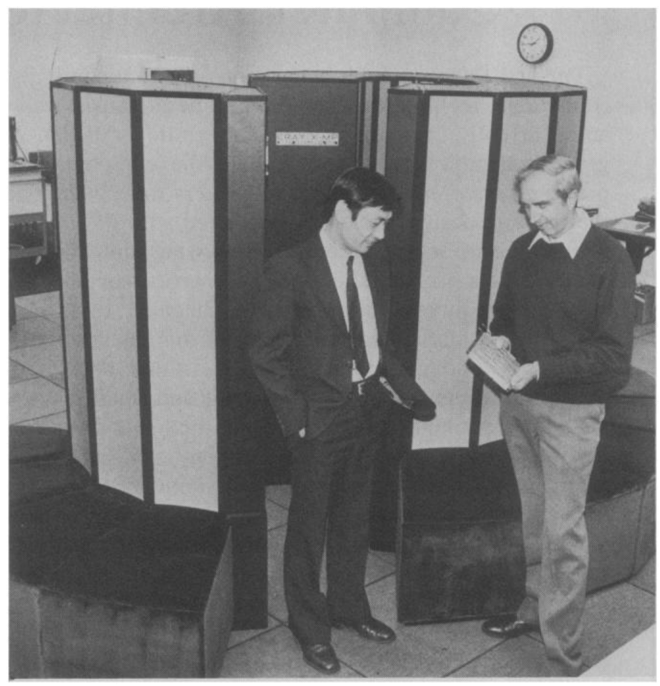

Cray Research vice president Steve Chen (left) and executive vice president Les Davis with a first version of company's soon-to-be marketed X-MP series supercomputer.

**Report on Large Scale Computing in Science and Engineering* (National Science Foundation, Washington, D.C., 1982).

Carnegie-Mellon, Massachusetts Institute of Technology, and Stanford.

Recently, DARPA seized the initiative by announcing its own supercomputer program. The agency director, Robert S. Cooper, told a House Appropriations panel on 16 March that DARPA planned to spend an additional $50 million in fiscal year (FY) 1984 and $95 million in FY 1985 "to develop the new generation of supercomputers and to explore this domain with a goal of achieving [enhanced] capabilities in our defense systems of the 1990's." Defense Department sources say more details of the supercomputer program will be furnished in a forthcoming DARPA position paper.

According to NSF estimates, a total of about $200 million is being requested in the FY 1984 budget for computer research in all federal agencies. DARPA would account for about half of total government expenditures in the field next year, with its new $50 million supercomputer program accounting for the bulk of the increase beyond the $137 million-plus being spent for all computer research this year.

NSF has a continuing interest in large-scale computing because of its importance to scientific and engineering research and has been a longtime federal patron of R & D in computer science and engineering. Precise figures on NSF support relevant to supercomputers are difficult to obtain because such projects are scattered through the budget, but a report to the National Science Board (NSB) recently put projected agency spending on computer research in FY 1984 at $51.5 million, up $8.5 million over this year. The NSB is scheduled to discuss possible expansion of computer related R & D at its May meeting.

DOE leads federal agencies as a proprietor of supercomputers. With an interest in large-scale computing prompted by its responsibilities for nuclear power and weapons research, DOE owns about half the more than 25 Cray and Cyber 205 supercomputers the government has bought. DOE operates a magnetic fusion energy computing network open to agency contractors and grantees. DOE also runs a relatively modest research program focused on parallel processing architecture. Funding of about $4 million is in next year's budget, double what was spent in 1980.

NASA also has a vested interest in large-scale computing. Because of its responsibility for aircraft design, a main concern is to be a world leader in computational fluid dynamics. This is translated into having the best supercomputers. The agency's aim now is to acquire what it calls a numerical aeronautical simulator, actually a network of computers based on supercomputers. NASA thinking now is that it will have congressional and Office of Management and Budget approval to spend an initial $20 million next year when the promised Cray-2 machine becomes available. NASA ultimately plans to spend a total of about $100 million on the network.

NASA's attitude, according to Jack Kerrebrock, the agency's associate director for aeronautics and space technology, is to "be a friendly customer for the best technology the companies can deliver." Kerrebrock says there is a good reason for this outlook. "Intellectual resources are very precious. Only a few people can conceive and design these machines. We don't want to use unique resources to respond to government specification," but would rather encourage companies "to develop a product line."

There is no single government view. The prevailing attitude at DOE, for example, differs from that at NASA. DOE has followed a practice, started by its forerunner agency, the Atomic Energy Commission, of encouraging supercomputer development by issuing specifications for advanced machines to meet agency requirements, then buying initial models of the resulting machines and, in many cases, developing the software for the new machines. DOE's director for energy research, Alvin W. Trivelpiece, sees the need for federal agencies to

Computer Architectonics

The first four computer generations are commonly distinguished by their constituent technologies—vacuum tubes, transistors, integrated circuits and, currently, very large scale integration (VLSI). Central to the fifth generation concept is a break with the computer architecture that has prevailed since the first digital computers were built. The basic elements of the existing design, which is associated with the mathematician John von Neumann, are a single central processing unit linked to a memory. The communication channel between the processor and the memory restricts performance by requiring operations, in effect, to be performed one by one, though at rapidly increasing speeds as computers have been improved. The consensus among computer experts is that despite further gains to be obtained through use of new materials and smaller machine geometry, von Neumann machines are approaching the limits of performance.

Now, advances in VLSI technology are raising prospects of breakthroughs in computing power and flexibility. Of particular interest are new computer architectures employing processors in parallel. A major goal is what is called "massively parallel" configuration with a thousand or more processors linked to literally thousands of memory units.

Development of new architectures is an aim of the Japanese Super Speed Computer Project, but it is the promise of innovation in architectures that has encouraged them to bank on incorporating sophisticated artificial intelligence (AI) functions in the fifth generation machines.

Research on AI, pioneered in the United States and Europe, ranges from exploratory studies such as computer modeling of human thought processes to more narrowly focused attempts at problem-solving. A pragmatic line of research has resulted in development of so-called "expert systems," which make specialized knowledge available through computers to perform demanding tasks such as medical diagnosis, chemical analysis, DNA sequencing, or identification work of various kinds for the military. Creation of so-called knowledge-based systems is at the heart of the Japanese fifth generation effort.

Problems of terminology are cropping up in discussions of supercomputers. Supercomputers have been defined as the machines offering the greatest speed and largest memory at any given time. Successive generations of supercomputers were progressively more powerful "number crunchers," able to do numerical calculations at rapidly increasing speeds. Prospects for development of machines with highly sophisticated AI capacities and of so-called special-purpose machines operating at very high speeds to perform specialized tasks means that advanced machines in future will be superior in different ways.—**J.W.**

continue to play such a role if only because the market for supercomputers is small and uncertain, which discourages U.S. companies from assuming the costs and risks of supercomputer development. Trivelpiece says his personal view is that it is axiomatic that the United States should not relinquish the leadership in supercomputers without a stiff fight. He sees a possible government role, for example, in guaranteeing government purchases of supercomputers to "create enough of a market to make it worthwhile for the companies."

Some observers note that the Japanese have bought only two American supercomputers, but expect that domestic demand for the big machines will suddenly spurt when Japanese supercomputers become available.

The main forum for the discussion of federal policy on supercomputers is a committee of the interagency Federal Coordinating Council for Science and Technology (FCCST), which was formed last year. The committee is overseen by the White House Office of Science and Technology Policy. Charged with looking at the general problem of supercomputers, the committee at its second meeting in early April formed working groups to consider ways to expand the user community for supercomputers and to assess research in the field government-wide. The committee's deliberations are at an early stage, but the record of FCCST (pronounced fixit) is one of modest success in coordinating federal activities, not of producing dynamic policy initiatives. A strong influence on the current federal outlook is thought to be the Administration's philosophic preference for limiting the government role in the economic sphere and relying on the working of the market. Also a factor is skepticism about the capacity of the U.S. bureaucracy to orchestrate a government-industry partnership in comparison to Japan's mighty MITI (Ministry of International Trade and Industry).

U.S. industry has made its own collective response to the Japanese challenge. The major direct reaction has been formation of the Microelectronic and Computer Technology Research Corporation (MCC), which is characterized as a research cooperative. MCC has 16 corporate sponsors and plans a central staff of 250 scientists and engineers. MCC's work is not aimed exclusively at supercomputer development, but many of its projects, such as those on artificial intelligence, parallel processing architectures, and intelligent systems, are directly relevant. MCC's budget is expected to rise to $50 to $80 million a year. Its director is former CIA and National Security Agency official Admiral Bobby R. Inman, who is being counted on to steer the organization through the shoals of the antitrust laws which have restricted industry-wide planning and cooperation. Industry has also set up a Semiconductor Research Cooperative (*Science*, 6 August 1982, p. 511) which is expected to have a budget of as much as $35 million a year by 1985 and should also contribute to semiconductor technology.

At this point, American investment in R & D relevant to supercomputer development appears to exceed Japan's, particularly if the big research programs of IBM and other major computer manufacturers are taken into account. But a deep impression has been made here by the capacity of Japanese government and industry to work in concert to achieve ambitious goals through planning, mobilization of resources, and coordination of effort. Japan's ability to overtake the United States and surpass it in targeted areas of technology was demonstrated most recently with semiconductors.

The most serious obstacles confronting the United States in competing head to head with the Japanese on supercomputers appear to be those identified during the Carter Administration inquest on declining U.S. innovation and productivity in U.S. industry. The workings of U.S. antitrust legislation and of patent, tax, and trade policies are seen as handicaps in the contest with Japan. Current American management and financial practices are also blamed for sapping U.S. competitive vigor. And adversarial attitudes in most relations involving government, business, and labor are compared unfavorably with the cooperative attitudes said to be dominant in Japan.

The Japanese, however, hardly face a clear track to preeminence in supercomputers. Cray Research and Control Data Corporation, which makes the Cyber 205, are not retiring from the field. A spokesman for Cray, for example, says that the performance announced for the Fujitsu and Hitachi machines exceed that of the Cray-1, but are "not up to the performance" of the Cray-X-MP available later this year. He said the X-MP will be two to five times more powerful than the Cray-1 and that the Cray-2 due in 2 years will be two or three times more powerful than the X-MP. As for the

Japanese Designs

Japan's aspirations for the Fifth Generation Computer Project are reflected in the differing organization of its two national computer projects. The Super Speed Computer Project is based on a partnership involving Japan's Ministry of International Trade and Industry (MITI) and six major computer firms—Fujitsu, Hitachi, Mitsubishi, Nippon Electric Company, Oki, and Toshiba. The project was launched last year and is scheduled to run through 1989. Total funding by MITI and the participating firms is projected at $200 million. As in past collaborations, R & D projects will be distributed among company and government laboratories.

The fifth generation project, while it is also based on a MITI-sponsored partnership, departs significantly in organization from earlier Japanese catch-up programs. The centerpiece of the project is a new laboratory, the Institute for New Generation Computer Technology (ICOT) presided over by a director, Kazuhiro Fuchi, who has been granted an unusual degree of responsibility and managerial latitude for Japan, where in most things, even leading-edge technology, consensus rules.

ICOT has started with a staff of 40 young computer scientists and engineers selected by Fuchi from the participating companies. These staff members will return to their companies after a couple of years perhaps as converts who will be effective missionaries for an ICOT style of innovative R & D.*

The Fifth Generation Computer Project also started in 1982. It is scheduled to run for 10 years; funding has been put at some $500 million from government and industry sources over the decade. But observers say that staff salaries and some other costs are being assumed by the companies and, therefore, the budget may be understated.—**J.W.**

*The aims and organization of the fifth generation project and its AI background are described in a forthcoming book, *The Fifth Generation*, by Edward A. Feigenbaum and Pamela McCorduck, published by Addison-Wesley.

future, Japanese progress "got our attention," he said, but "Cray Research is dedicated to one proposition: to continue to provide the fastest, most powerful machines in the world."

For the Japanese also, formidable problems in computer architecture and component technology remain to be conquered and they have yet to shatter the stereotype that they are talented copycats. There are also some doubts that ingrained Japanese caution in management and reliance on consensus in decision-making is well suited to innovation. And Japanese universities at the moment are regarded as incapable of turning out the cadre of computer scientists and engineers required to fulfill Japanese aspirations.

Nevertheless, is it possible that the United States and Europe could continue to excel in innovative work and the Japanese continue to read the journals and visit the Western labs and use their superb development and production talents to build a dominant role in supercomputers?

One federal scientist in the thick of a supercomputer project concedes, "It's a real concern. We worry about it an awful lot. It's the kind of technology transfer we'd like not to have." He said it is clear that a "better interchange between research and industry" is needed in the United States. "But if we were to do something that would tend to discourage the free exchange between researchers we would destroy the very synergy we sought to create. We have a great technological society but we Americans have got to emphasize not tripping over our own feet."

The dilemma over supercomputer development is not unique. Rather, it puts in acute terms the question of the U.S. rivalry with Japan in high technology and of what the U.S. government should or shouldn't do about it.—**John Walsh**

Japanese Borrow Plan from U.S.

Japan's plans for its fifth generation computer systems are audacious and daring. But, with little experience in the areas they wish to pursue, how did the Japanese formulate them? With the help of American computer scientists, apparently. Researchers from Massachusetts Institute of Technology (MIT) in particular but also from Stanford and Carnegie-Mellon University seem to have had an overriding influence on the plans.

Michael Dertouzos, head of the Laboratory for Computer Science at MIT says he "panicked" when he first saw the Japanese plans, feeling it was somehow not right for the Japanese to copy MIT ideas so blatantly. Now, he says, he has changed his mind. "More power to them," he says. He only wishes American industry would listen to the university scientists so attentively.

What the Japanese did was to invite eminent American computer scientists to visit and lecture in their country and describe what they thought would be important projects to pursue. Among the MIT researchers invited to visit Japan were Jack Dennis and Gerald Sussman. Dennis suggested that the Japanese build data flow machines, an MIT invention. The data flow machines proposed by the Japanese are, says Dertouzos, "unmistakable" in their MIT origins.

Sussman gave the Japanese his ideas on artificial intelligence. He was invited to visit Japan for 3 days in November 1979. Although he did not know the purpose of his visit before he arrived, it turned out that he was to speak before Japanese scientists planning the fifth generation project. "For 8 or 9 hours a day I was flaming forth," he recalls. "All their distinguished scientists were there. They'd clearly read everything I had ever written and had Japanese translations of my papers in front of them. They were very smart, very well prepared, and they kept their mouths shut."

Sussman's ideas were incorporated almost verbatim in the Japanese proposal. Even one of his slides and the diagrams he drew on the blackboard appeared. Sussman recalls that he looked at the Japanese report on their long-range plans with surprise, seeing his proposals and saying to himself "Gee, that looks like a Sussmanism." In fact, he says, the original Japanese report "is so close to an MIT research proposal that it's frightening."

Edward Feigenbaum of Stanford University says his visit to Japan resulted in the Japanese plans to emphasize knowledge engineering or expert systems. Feigenbaum spent 12 weeks in Japan in 1979 during which time he gave lectures at major universities and industrial laboratories. "In all these places I gave lectures on applied artificial intelligence at Stanford. This is work on expert systems. The result was to make an extremely convincing case to the Japanese that this was a hot area to invest skill and energy," Feigenbaum recalls.

As a result, Kazuhiro Fuchi, who was then head of the information science department at Japan's Electrotechnical Laboratory and who is now director of the Institute for New Generation Computer Technology (ICOT), the central laboratory for Japan's fifth generation project, convinced the Ministry of International Trade and Industry (MITI) that it should include expert systems in its long-range plans. "The credibility was supplied by my lectures," Feigenbaum says. "Some Japanese referred to me as the Father of Knowledge Engineering in Japan."

Feigenbaum explains that using his statements to lend prestige to a project is firmly in the Japanese cultural tradition. "The Japanese have a saying, 'The nail that sticks its head up is the one that gets hit.' If they can say, 'Feigenbaum had success with this idea,' it's my nail that's sticking up."

In Feigenbaum's new book, *The Fifth Generation: Artificial Intelligence and Japan's Computer Challenge to the World*, written with Pamela McCorduck, he tells of asking Fuchi whether Japan could really build expert systems with so little experience. He quotes Fuchi as replying, "Metaphorically speaking, if your countries are like adults, then Japan may be likened to a baby, but in my own mind Japan is actually closer to boyhood. It may seem funny for me to talk about how a boy should behave, but boys must learn from adults and listen to them and respect their opinions."

Asked how he feels about Japan's adoption of his ideas, Feigenbaum says, "I have mixed feelings. I feel very good about my ideas being used but I also worry about the health of our own information processing industry. I would much prefer ideas to go from Stanford to American industry."

—**Gina Kolata**