Supercomputers: who'll be the fustest with the fastest?

ALEXANDER/THE SAN FRANCISCO EXAMINER

"We've discovered the most primitive tribe in the world . . . They have no word for computer."

In the race to develop faster, cheaper, more intelligent "supercomputers," the US seems to be taking the approach of the tortoise. According to some watchers of the race, the hares are Japan and Europe. Though the US computer industry dominates the field with 85 of today's 87 most powerful machines in operation in the noncommunist world, many members of Congress and the computer-science community fear the Aesopean outcome may not prevail in the supercomputer race.

Thus, at the outset of hearings on supercomputers before the House Science and Technology Committee last November, its chairman, Don Fuqua (D-Fla.), expressed "the desire to be number 1" and argued that "Japan is threatening our entire computer industry." A few weeks earlier, Michael L. Dertouzoz, director of the MIT Computer Laboratory, characterized the threat in *Newsweek* as another Pearl Harbor waiting to happen. Meanwhile, Edward Feigenbaum, co-author (with Pamela McCorduck) of *The Fifth Generation* (Addison-Wesley, 1983), cautions that the US possesses the capability but lacks the resolve to beat Japan to a new breed of supercomputers. This assessment is endorsed by a National Science Foundation panel headed by Peter D. Lax of New York University, which reported after a comprehensive study of the situation that "US leadership in supercomputers is crucial for the advancement of science and technology, and therefore for economic and national security," then warned of "alarming" evidence that "US dominance of the supercomputer market may soon be a thing of the past" (PHYSICS TODAY, September, page 42).

What lurks behind such worries is not that Japan or some other country will wrest control of large-scale computing from the US in the next few years, but the shape of things to come. Pre-eminence in the computer industry is certain to have wide implications. Computational dynamics is likely to transform entire industries and stimulate new technologies. Some of the computationally most intensive problems involving such disciplines as aerodynamics, seismology, meteorology and atomic, nuclear and plasma physics may be solved by the prodigious power and speed of information-processing machines. One of the great virtues of supercomputers is their ability to replace laboratory experiments, so that they can be used, for instance, by pharmaceutical companies to develop new drugs with molecules derived from quantum theory, or they can create a graphic three-dimensional model of a fusion reactor in operation.

Not since Galileo. After the first meeting last January of the NSF's 15-member Advisory Committee for Advanced Scientific Computing, its chairman, Neal F. Lane of Rice University, issued a manifesto that most participants accepted. It reads: "Science is undergoing a structural transition from two broad methodologies to three—namely from experimental and theoretical science to include the additional category of computational and information science. A comparable example of such change occurred with the development of systematic experimental science at the time of Galileo.

"The current transition affects every discipline of science and engineering and every university engaged in scientific research. We expect that the new opportunities for scientific inquiry and industrial applications of science made possible by large-scale computational support will feed back into total scientific activity—in theory, experiment and computation. Industry is becoming increasingly dependent on large-scale computation to help design products that are competitive in world markets; its need for breakthroughs in research and for trained manpower should help swell scientific activity."

From this perspective, it is not surprising that Japan and Europe have decided for their own futures to become rivals of the US in developing supercomputers. Japan was first to come to

0031-9228 / 84 / 0500 61-04 / $01.00

grips with the stakes. In October 1981, it revealed plans to produce supercomputers 100 to 1000 times faster than today's Cray-1 and X-MP or Control Data's Cyber-205 before the end of the decade. Shortly thereafter, Japan would provide a new breed of "fifth-generation" systems. Both projects embrace six Japanese firms (Fujitsu, Hitachi, Nippon Electric Corp., Mitsubishi, Toshiba and Oki Electric) along with MITI, the government trade and technology organization that plans and funds industrial policies. Together, the companies and MITI operate in a collaborative, coordinated way. The first project, the $100 million National Superspeed Computer development, is already coming up to speed. Signs of this abound. In recent months, even though not formally part of the project, a Fujitsu VP-100 has been delivered to Nagoya University and a Hitachi S-810/20 to Tokyo University. Nippon Electric's SX-1, reputed to be superior to these, should be delivered next year, if all goes according to plan.

The other project, the $500 million fifth-generation venture, aims more ambitiously at developing machines using an array of high-speed logic and memory devices (Josephson junctions, high-electron mobility transistors, gallium arsenide microcircuits) and parallel architectures. The machines are intended not only to process mathematical data but also to provide human-like reasoning or artificial intelligence, to recognize, manipulate and draw inferences from non-numerical symbols—words, sentences and visual or aural patterns. As such it is a radical break from conventional computers based on concepts developed by two brilliant mathematicians, Alan Turing at Cambridge and the National Physical Laboratory (England) and John von Neumann at the Institute for Advanced Study, during and immediately after World War II.

Generational bottleneck. Since the first digital computers, the machines have followed von Neumann's general design, consisting of a single central processing unit linked to a memory by a communications channel, which, it turns out, restricts performance by requiring operations to run, in effect, one by one. Computer designers generally agree that despite further gains to be had by using new materials and smaller machine geometry, von Neumann computers have remained virtually unchanged for 30 years and are reaching the limits of performance. The communications connection has come to be known as the "von Neumann bottleneck." Efforts to eliminate the bottleneck have been made in all four generations of computers. The generations are customarily identified by the basic inner components of computers:

- electronic vacuum tubes
- transistors
- integrated circuits
- or very large-scale integrated circuits.

The fifth generation, says Feigenbaum, "will stand apart not only because of its technology but also because it is conceptually and functionally different from the first four generations the world is familiar with. The new machines will be known as knowledge information processing systems—or KIPS."

Non-von concepts. "We think we can diverge from von Neumann principles of machine logic in the fifth generation," says Kazuhiro Fuchi, director of the Research Center of the Institute for New Generation Computer Technology (or ICOT, when translated from the Japanese). "We will replace the established culture of today's computer world with a new culture based on natural logic, with new architecture, new software and new applications." In achieving this, Fuchi says, Japan's computer scientists and engineers will have to develop understanding of human problem-solving processes so that sophisticated artificial-intelligence functions can be incorporated into fifth-generation machines. One idea for the "non-von" computer that ICOT is pursuing would eliminate the internal bottleneck by intermingling a massive array of processors and memory units throughout the machine.

Europe also plans to seize the moment of technological transition to try to overcome US computer dominance. In Britain, some $530 million is committed by government and industry equally over the next five years to artificial intelligence, VSLI research and software engineering. The project is a collaborative effort by government, industry and universities. The West German government is contributing $4 million per year to universities for R&D in parallel processing, and next year France has agreed to increase its spending for more computer specialists at its national research centers.

What's more, the ten-member European Economic Community, after wrangling nearly nine months over political, organizational and financial problems associated with an advanced computer project, finally agreed in March to put together a five-year, $1.25-billion R&D effort. The work on VSLI technology, parallel hardware and AI software will be done by universities and 12 private computer and electronics companies, including Siemens, Philips and Britain's General Electric, which are expected to match the funding by governments. The project bears the official title of European Strategic Program for Research in Information Technology, a mouthful that has been given the felicitous acronym of ESPRIT.

No 'crash' program. The projects undertaken by Japan and Europe have sent a chill through many US corporate boardrooms and university research centers. Like Feigenbaum and Lax, Jack Worlton of Los Alamos believes that without stronger Federal support for computer research at universities and national labs, "the American supercomputer industry will be overwhelmed by the Japanese and possibly Europeans in the long term." But in testimony to Congress, President Reagan's science adviser, George A. Keyworth II, dismisses such dire predictions. The government, he insists, remains committed to "maintaining our historic lead in developing successive generations of ever-faster, general-purpose computing machines. It intends to continue to maintain a climate in which US supercomputer companies will have the opportunity to remain in the lead," without resorting to a "crash" program or Manhattan Project to respond to the new challenges from abroad. Reflecting the Reagan Administration philosophy, Keyworth declares that the government will not intervene in the industry's "own choice and judgment about the technologies and markets for advanced computers."

Keyworth argues that the US computer industry is moving swiftly by increasing its R&D funds for supercomputers and that the proposed National Innovation and Productivity Act, which is wending its way through Congress, will encourage such collaborative industrial efforts as the Microelectronics and Computer Technology Corp. (which brings together 14 companies including Texas Instruments, Motorola and Digital Equipment) and Semiconductor Research Corp. (made up of 23 firms, including Hewlett-Packard and National Semiconductor). What's more, as an example of cooperative research, the Department of Defense and some 20 companies have established the Center for Integrated Systems at Stanford, where work is going on in robotics, AI and microelectronics. Other cooperative activities are under way at Cornell, Carnegie-Mellon, MIT and Rensselaer Polytechnic Institute.

HEPs and PUPS. The government will continue to buy supercomputers to meet its own needs at the national labs for nuclear weapons design, say, and weather forecasting and space exploration. Keyworth noted that prototype and pre-production models will be supported for development in government and university labs or installed in those places when built by commercial computer companies. Thus, at Los Alamos, where Keyworth once led the laser fusion program, a processor called HEP

Faster than a speeding bullet or its shock wave

How fast is a supercomputer? Faster than a "speeding bullet," the headlong velocity of Superman in pursuit of the world's villains. Faster than a Goose Gossage fast ball, with the wind at his back. Faster, even, than a millisecond pulsar.

In the realm of supercomputers, speed is measured in flops—floating point operations, or calculations, per second. Today's Class 6 supercomputers are rated in megaflops—Cray-1, built by Cray Research Inc., at about 100 megaflops, Cray X-MP some 400 megaflops (roughly 210 million calculations per each of two CPUs), Cyber-205 of Control Data Corp. at 200 megaflops (with two accelerated pipelines and 400 megaflops with four pipelines). Thus, in the blink of an eye or the twinkle of a star, a supercomputer can start cranking through Navier–Stokes equations, the mathematical puzzles that can stump a wind tunnel full of aeronautical engineers for months.

NASA and the Department of Energy's national laboratories say they need to solve problems that require computer capacities of 10 gigaflops and memories of at least 256 million words. The Cray-2, which is undergoing its final tests and should be available early next year, is most likely to reach 1 gigaflop and contain addressing space for 256 million 64-bit words. Currently, the Cray-1 memory is 4 million 64-bit words, or up to 32 million bytes, and the X-MP's is 4 million words.

Japan introduced two new Class 6 supercomputers this year, based on a different design philosophy from US machines. The VP-100 of Fujitsu Ltd. and S-810/20 of Hitachi Ltd. both accentuate vector performance over scalar and, in consequence, caused computer analysts to misjudge their speeds. Originally, they were said to run at around 500 megaflops at peak rate. Sustained performance will be much lower, of course. The Japanese machines range from 5 to 500 megaflops. Fast, but not out of the US ballpark. —IG

(for Heterogenous Element Processor) produced by a small new firm, Denelcor Inc., has eight simultaneous streams of information flowing through its architecture. Los Alamos is developing its own experimental processor known as PUPS (for Parallel Micro-Processor System, with a U replacing the μ for micro), which is intended to extend the computational capability of a Cray-1 by another 10%.

Other government programs are aimed at increasing the number and quality of computer scientists and engineers and enabling university and company researchers to gain greater access to supercomputers. Last year, says Keyworth, the US produced only 225 PhDs in computer science and engineering, while in the field of computational mathematics fewer than 30 received PhDs. Nearly all were snapped up by computer and electronics firms. NSF reckons that US universities are short about 200 faculty to teach and do research. "We need to do something quickly about that situation," declares Keyworth. "But the problem doesn't lend itself to a quick fix."

Keyworth's agenda for supercomputers is similar to that of the Lax panel, an NSF working group that issued a report last year entitled "A National Computing Environment for Academic Research," and his own ad hoc Committee of the Federal Coordinating Council on Science, Engineering and Technology or FCCSET (pronounced "fixit"). None recommended a national program such as the ones established by Japan or Britain, but all urged that the US government do such things as support R&D essential to the design and development of new supercomputer systems, increase training in computer sciences and engineering and improve access to scientific computers.

Federal actions. According to NSF estimates, a total of $350 million in all Federal agencies is requested for computer research and training in the fiscal 1985 budget—nearly double the amount obligated last year. The Defense Advanced Research Projects Agency, which has inherited a long history of support for computers, going back to the ENIAC at the end of World War II, accounts for nearly half the total government expenditures in the computer field, much of it for the VHSIC (for very high-speed integrated circuits) program and for President Reagan's Strategic Defense Initiative (better known as "Star Wars"), which requires ultra-fast supercomputers to respond to a ballistic missile attack, when seconds count. At a supercomputer conference in Los Alamos last August, DARPA director Robert S. Cooper announced that his agency would put at least $500 million into academic and industrial research on advanced computers in the next five years. Hearing that, his boss, Richard DeLauer, Under Secretary for Defense for Research and Engineering, interjected: "What is needed for R&D in that field during the period is probably $3 billion."

NASA has a vested interest in large-scale computing because of its responsibility in aeronautical design and the needs of the space program. Its labs at Ames, Goddard, Langley and Lewis have a total of two Cray-1s, two Cybers and a CDC-7600, which are available to NASA grant-holders without charge.

The Department of Energy leads all the Federal agencies as a proprietor of supercomputers, with more than a dozen Crays and Cybers stationed at the national laboratories. While most of these machines are for research in nuclear energy and weapons work, some are for wider use. Of the nine supercomputers at Los Alamos, a Cray-1, and three CDC-7600s are "outside the fence," meaning they are accessible to government agencies, other labs and non-profit institutions such as universities, when these share common programmatic interests, at a cost of $636 per prime CPU hour. In addition, since 1975, Livermore has been operating the National Magnetic Fusion Energy Computing Center for contractors in DOE's fusion program. The center is used free of charge by some 2000 scientists and engineers in 50 labs, universities and companies who access its two Cray-1s and a 7680 through a network linked by satellite and land lines. A new Class 7 computer, likely to be a Cray-2, is budgeted for installation at the end of the year, though the schedule is apt to slip to 1985. DOE has begun to increase the number of users to include researchers under contract with it in high-energy and nuclear physics, basic energy sciences and health and environmental programs. Before the additional users can be accommodated, however, more capacity will be required. Accordingly, DOE plans to lease another Class 7 computer in fiscal 1985, when the budget earmarks nearly $24 million for the Livermore fusion computing center.

Precise figures on NSF support of computer research and researcher access are difficult to establish because such programs are scattered through the budget, but a report to the National Science Board put the expenditures in fiscal 1984 at $51.5 million, up to $8.5 million above 1983. NSF's new budget calls for more than $100 million for various computer programs. Of this $39.4 million, some 16% more than 1984's $33.9 million, would go for computer research, which covers theoretical studies, artificial intelligence, VLSI designs, systems architectures and other basic computer sciences conducted mainly at universities. Last year NSF supported 32% of all computer research in universities, including virtually 90% of theoretical science and 50% of software R&D.

When NSF first drew up its budget request for fiscal 1985, it sought $47 million to expand access to advanced computer facilities for those scientists and engineers not eligible to tap into research supercomputers under control of other agencies—namely DOD, DOE and NASA. The Office of Management and Budget lopped this item to $20 million before the budget went to press. In Congress, however, NSF was voted another $40 million for computer access, more than restoring the OMB cut. NSF has not produced a fully developed plan for access to supercomputers, but even the $60 million allocated for the coming year, which represents an enormous jump over the $6 million available this year, pales against the rough estimate made in the NSF staff report in 1983 that more than $400 million would be needed over the next three years to deal adequately with the problem.

The staff report concluded that while there is an immediate need to make supercomputers more available to academic scientists and engineers, "the attitude toward computers for research marks a generation gap in the community. . . . Older scientists may use, or have their graduate students use, computers for a variety of measurement and analysis tasks but consider the role and importance of computers to be limited. Younger scientists, who grew up with computers as part of their culture, view them differently. . . . Science has passed a watershed in using computers for research. Computers are no longer just tools for measurement and analysis, but large computers in particular have become the means for making new discoveries."

In sum, then, while there is no clear government scheme to coordinate a coherent national supercomputer program, which would embrace research, development and access, it is becoming clear that the acute rivalry with Japan and Europe will force the administration and Congress to become more responsive to developments here and abroad. —IG

SSC cost and size perplex Congress

Sometimes the history of physical science resembles nothing so much as political science. For instance, in the 1950s and 1960s particle physicists had learned enough about political methods to establish alliances with congressmen and governors in their vigorous and often acerbic battles against each other for large and expensive high-energy accelerators. Professional politicians, for their part, often expressed their vexation with the combat over "big science." Thus, early in 1964, soon after becoming president, Lyndon Johnson, weary of the internecine warfare among physicists over the site of the 200-BeV proton accelerator (built later at Fermilab), wrote to Hubert Humphrey, then a senatorial champion of one of the factions: "I devoted more personal time to this problem than to any nondefense question that came up during the budget process."

Members of the House Subcommittee on Energy Development and Applications voiced somewhat similar dismay on 22 February as they listened to physicists explaining the need for a Superconducting Super Collider—the proposed crown jewel of today's big science. In his testimony before the subcommittee, Alvin W. Trivelpiece, director of energy research at DOE, said $20 million is requested in fiscal 1985, only $1 million more than is appropriated this year, to carry out exploratory evaluations of various concepts of the powerful accelerator, its magnets and cryogenics.

The work in 1985 will be based largely on the findings of a reference design study involving some 40 scientists and engineers from seven universities and national laboratories under the leadership of Maury Tigner of Cornell. The Tigner group has been pondering the principal problems of building and operating a high-luminosity proton–proton collider with 10 to 20 TeV in each beam—namely, the type and strength of superconducting magnets, the cryogenic requirements and conventional facilities, as well as the "credible" cost estimates for the project.

Skepticism. Representative William Carney (R–N.Y.) asked Tigner, "What is Congress to do?" Said Carney: "We're dealing with the largest scientific endeavor in our nation's history. . . . There is some skepticism by members of Congress as to whether we should go forward with the project." Claudine Schneider (D–R.I.), only ten minutes before, reminded Trivelpeice and the physicists who testified, including Tigner, Leon Lederman of Fermilab and Burton Richter of SLAC, that "when we are looking at a $180 billion deficit in 1985, [the SSC] will require an enormous amount of justification."

Carney claimed Congress was receiving "mixed signals." Scientists are advising the government to proceed with the SSC as rapidly as possible to wrest the leadership in high-energy physics from Western Europe, said Carney, but DOE cautions against moving too far too fast lest the project "lock in" one approach that later proves to be too costly and unworkable. The Tigner study, which concludes with a report to the DOE in May, will compare three different concepts but not necessarily make a choice among them for the final design. The report will be discussed at a conference on particles and fields at Snowmass, Colorado, in late June, and used by Energy Secretary Donald P. Hodell to decide on the government's intentions for the SSC in preparing the DOE budget request for fiscal 1986. "Will $20 million be enough in fiscal year 1985," asked Carney, "if the decision is to go forward with the project?"

"Well, I'm not terribly happy with it," replied Tigner. "We are trying very hard now to make a convincing and detailed argument about what is required to go forward, and we are very hopeful that we will be able to convince the DOE to support a somewhat broadened scope for the activities."

Direction? Though he is an advocate of the behemoth collider, Carney admitted, "I've become confused." Last October, he recalled, DOE's High Energy Physics Advisory Panel testified before the same subcommittee that Congress should support the SSC, though the HEPAP members allowed that the design, location and cost of the machine were still uncertain. "What decision do we make in Congress? Just wait another year, spend another $20 million, with no direction? . . . We might be able to reduce the deficit," Carney continued, by not funding the project, "but that would be to the detriment of high-energy physics in our country."

Cost, of course, is only one consideration in congressional decisions. Almost as important is the location of a particular project or facility. Accordingly, Representative Schneider asked: "Are there any thoughts as to where the SSC would be placed?" "There are a lot of thoughts about where it might be located," Trivelpiece responded. "Hardly a week goes by in which I do not get an offer to have it in some particular location." Representative Paul Simon (D–Ill.) wanted Trivelpiece to know that he had considered a site "very carefully, very objectively" and concluded, not surprisingly, "that Illinois is a natural location for the proposed machine." Carney noted that the SSC "just would not fit" either in Schneider's Rhode Island or in his district on New York's Long Island, since preliminary plans suggest that the accelerator would be some 12 to 32 miles in diameter.

One of the most far-reaching suggestions for a site, Trivelpiece interjected, "was to arrange the SSC in such a way that its beams, which can be steered with magnets rather effectively, will be directed to pass through the physics departments of about 20 different universities on the eastern seaboard. In that way, a support base could be assured." To this, Carney said with a chuckle, "If that is the design plan, high-energy physics would be operating like the Pentagon. Put it in 218 districts of Congress and any program will be funded." More seriously, however, Carney added, "I am frightened that in the budget process, we will say yes to the funds asked for now and then not have money to go forward agressively, pushing the SSC off another year. Delays cost money and also may cost us the support of Congress." —IG