Reaching for the Supercomputing Moon

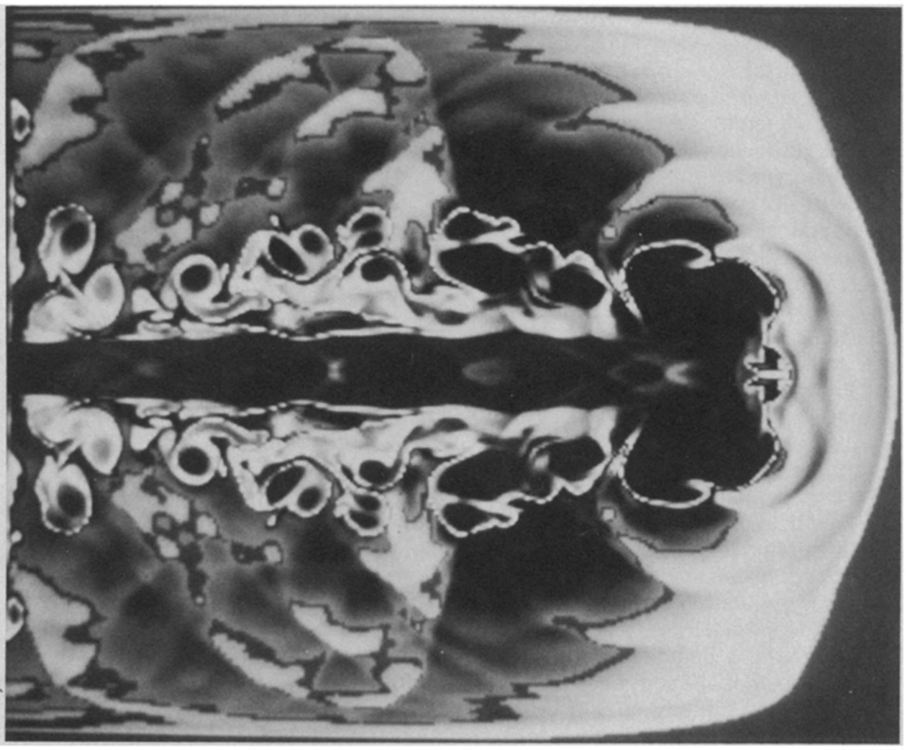

ETA Systems

A supercomputer can be used to model a gas jet traveling at supersonic speeds through a denser gas. An ETA^{10} supercomputer created this image in 15 minutes; a fast personal computer — if it had enough memory — would have taken 10 months.

By IVARS PETERSON

Putting a man on the moon was a dramatic challenge that captured the public's imagination. For more than a decade, this vision drove a massive, well-funded program that brought together government, industry and academia in a concerted effort to achieve the goal. That effort now provides a tempting model for a number of government officials, computer industry representatives and researchers who work with supercomputers. They would like to see a similar national program aimed at vastly increasing the capabilities of today's fastest computers.

"This is an extremely important technology," says Bill Buzbee of the National Center for Atmospheric Research in Boulder, Colo. "Supercomputing is so critical to our future that I'm reluctant to leave it to charity." He suggests that focusing on specific targets — such as building a computer capable of doing a trillion operations per second — may now be the best course to take.

Buzbee was one of roughly two dozen speakers commenting on the state of high-performance computing in the United States at a recent meeting before staff members of the House Subcommittee on Science, Research and Technology, which is preparing to hold formal hearings on the topic later this spring. The subcommittee's deliberations are likely to have a strong influence on the future direction of scientific computing in the United States.

This congressional activity follows in the wake of two significant reports: one issued late last year by the Society for Industrial and Applied Mathematics (SN: 11/21/87, p. 335) and a more recent study produced by a committee of the Federal Coordinating Council for Science, Engineering, and Technology (FCCSET), which oversees agency research programs. Both reports stress the importance of maintaining a lead in high-performance computing to remain economically competitive, especially with Japan. Both call for increases of at least $1.5 billion over the next five years in federal spending for programs related to supercomputing

There are plenty of scientific problems that require greatly increased computing power (SN: 9/29/84, p.200). The FCCSET report includes the following "grand" challenges in its list: computing the way a fluid flows over a large, complicated surface in order to improve aircraft design or to build quieter submarines; designing new materials by starting with a list of desirable properties and the fundamental characteristics of electrons in matter; and recognizing spoken words or identifying visual objects quickly and accurately. "Many of these," the report states, "could be considerably advanced by the use of computer systems capable of trillions of operations per second."

But are these challenges grand enough to catch the public eye and to prompt the government to loosen its purse strings? Enrico Clementi of the IBM Corp. in Kingston, N.Y., complains that the FCCSET list is old and uninspiring. The same challenges appeared in lists 20 years ago, he says. "We need new words, better examples." Moreover, the FCCSET report fails to establish specific goals and priorities. "I would like to see a bit more urgency," he says. "We should have a precise timetable."

One danger of "a moonshot approach to computation," says Harold P. Hanson, executive director of the House Committee on Science, Space and Technology, is the possibility of a letdown after the goal is accomplished. That's particularly relevant if the goal is not carefully selected and simply involves something like reaching a certain operating speed.

In fact, many of the problems that currently plague supercomputer use and development aren't even related to computational speed: There aren't enough supercomputers to meet the demand. No single network effectively links the supercomputer centers already in operation. The development of computer programs and improved computing methods, or algorithms, lags far behind advances in computer architecture. Supercomputer users lack the sophisticated methods necessary for handling the vast quantities of data that a supercomputer can store and process.

For many, the key issue today is the absence of a national research network linking computers of all colors and stripes. According to the FCCSET report, current commercial and government-sponsored networks are largely uncoordinated, have insufficient capacity and don't ensure privacy. Too little effort has so far gone into developing protocols that allow different types and brands of computers to communicate conveniently.

"There's a lot of duplicated effort because the machines are independent," says geophysicist Larry Brown of Cornell University in Ithaca, N.Y. He knows of several geophysics projects, all similar in scope and topic but going on at different supercomputer centers with little sharing of information. Another observer cites the example of a physics department in which individual professors have links to five different, separate networks.

Antony Jameson of Princeton (N.J.) University is concerned that too much attention is being focused on the fastest and largest supercomputers and not enough on much less expensive machines having near-supercomputer capabilities. "It's better to have pretty powerful computers all over the lot," he says, than to have a smaller number of large supercomputers that have to be shared by many individuals. Supercomputers should be reserved for problems that can't be done in any other way.

At present, more than 20 companies manufacture machines that perform as well as or better than the original Cray-1 supercomputer. Yet these "minisupercomputers" usually cost much less than $1 million. The trouble, says Jameson, is that researchers who would like to use such machines don't have any sources of funds for buying them. "The government bought the first few Crays," he says. In the same way, he suggests, it could help the fledgling minisupercomputer industry by providing funds to researchers who would like to use these computers.

The FCCSET report is primarily a government vision, says Paul G. Huray of the Office of Science and Technology Policy, which is coordinating government supercomputer efforts. "In terms of computing, we chose not to put a man on the moon." Instead, the emphasis is on coordinating and supporting the projects individual agencies propose for increasing computing power to meet specific needs. Senior officials in all of the government agencies concerned with large-scale, scientific computation are now working within their own departments and with one another to develop a plan for implementing the FCCSET report recommendations. So far, however, increased funding for supercomputing has failed to materialize.

Some critics of the FCCSET report contend that its recommendations offer little that is new. Edsger W. Dijkstra of the University of Texas at Austin, for example, argues that if it is true that the United States is losing its leadership in high-performance computing technology, then the report fails to analyze why that is happening and to place blame where it belongs. Instead, the report merely calls for a continuation of what was done in the past.

Nevertheless, making a national commitment and setting a timetable is important, says Hanson. "We need to establish a game plan for the nation." The hurdle is getting the funding needed to support such an effort. For supercomputer enthusiasts, that means finding an effective, eye-catching way to package the plethora of frequently conflicting goals and needs that permeate the present U.S. effort in supercomputing. ❑