

Parallel Paths

Hitching together thousands of simple processors produces a novel computer that does certain tasks faster than a supercomputer

By IVARS PETERSON

Making sense of data collected by satellite can be a formidable task. Just one sweep of a sensor across a 100-square-mile swath of the earth's surface produces a million-piece mosaic. Each tiny fragment of the image carries potentially useful information. Yet extracting that information, even from one picture, often takes days.

This is one of several image processing problems that led scientists at the NASA Goddard Space Flight Center in Greenbelt, Md., to commission the development of a unique computer known as the Massively Parallel Processor (MPP). Designed and built for \$6.7 million by the Goodyear Aerospace Corp. in Akron, Ohio, this experimental machine was delivered in May 1983.

One of its first tests involved analyzing data from the "thematic mapper" aboard Landsat 4 (SN:7/3/82,p.4; 8/7/82,p.84). By individually studying the million or so pixels making up a typical image, the MPP automatically works out whether each spot represents water or land, forest or field, stream or street — all in 20 seconds. A conventional computer that processes data one bit at a time would take hours to analyze the same picture and produce a similar classification scheme.

The MPP owes its speed to the unusual way in which the machine's parts are organized. Its network of 16,384 simple processors allows a computation to be divided up so that each processor performs the same operations on different pieces of data at the same time. The machine also has a "staging memory" that provides room for organizing data on their way into and out of the processing stage.

In December 1984, NASA invited researchers throughout the United States to suggest projects that would exploit the MPP's computational power. "We wanted to measure and document the advantages

and disadvantages of parallel processing," says Goddard's Milton Halem, space data and computing division chief. NASA selected 39 projects, and recently, at a symposium on "The Frontiers of Massively Parallel Computation," participating researchers described the results of their first year of work.

For Richard L. White of the Space Telescope Science Institute in Baltimore, the MPP provides a way to study the dynamics of galaxies containing hundreds of billions of stars. No computer is fast enough to track every star individually and to add up all the gravitational forces involved in such a system. Instead, White imagines these gigantic star assemblies as a special kind of fluid — galactic streams in which particles influ-


ence each other yet never collide.

White's computer simulations show cold, uniform distributions of stars evolving into lumpy patterns of hot, rapidly moving blobs or clusters with fine tendrils tightly wrapped around dense cores. "Slight changes in initial conditions give quite different final configurations," says White. "By experimenting with different initial conditions, we can learn about the factors that lead to the formation of galaxies."

For most of his computations, White found that the MPP is often faster than a Cray supercomputer. "A lot of features of the MPP are excellent for this type of application," he says. However, he discovered that computer programs for the MPP take longer to develop and are more difficult to modify later.



The MPP (center rear) at NASA Goddard with a VAX host computer (right rear).



One of a pair of synthetic aperture radar images (left) from which elevation data were extracted automatically by the MPP to produce a topographic map (right).

NASA Goddard

Computer scientist John A. Barnden of Indiana University in Bloomington is using the MPP to study how neural networks in the human brain may process complex information. His computer simulations attempt to answer the question: How may simple units that communicate in simple ways handle complicated concepts?

In his model, Barnden uses networks in the form of two-dimensional arrays — ideally suited for the MPP's checkerboard of processors. Each processor handles one grid square and keeps track of the square's "state" and "color," which represent properties of a small neural circuit. When new information is added and "reasoning" takes place, the squares change their properties as neighbors influence neighbors according to a set of rules. This theoretical model hints at how certain types of problem solving, such as identifying a shape, may occur in the brain.

A computer like the MPP is essential for this application, says Barnden. "The simulations involve too much computation to be feasible on conventional computers." However, the MPP requires all processors to work in parallel on the same task with direct contact only between immediate neighbors. That's a disadvantage for researchers looking at complex interactions between more widely separated squares.

The MPP and similar computers may also prove to be an inexpensive alternative to supercomputers for

running complicated simulations. Current computer models that describe the chemistry and transport of air pollutants over long distances, for instance, are now so detailed that they must be run on supercomputers. Even then, a Cray-1 takes about 100 minutes to come up with a prediction of what's likely to come down over the eastern United States one day after air pollutants are injected into the atmosphere. About 90 percent of that time is spent on atmospheric chemistry calculations.

Gregory R. Carmichael and Seog-Yeon Cho of the University of Iowa in Iowa City have now shown that the MPP can be used to speed up the chemical calculations, without making the model less realistic by simplifying its equations. In their program, each MPP processing element handles the chemistry calculations for a different "cell" in the grid covering the region being studied.

Preliminary tests show that the MPP executes these calculations several times faster than a Cyber-205 supercomputer. This means that researchers interested in problems like acid rain can run detailed models on the MPP or similar machines and get results within a reasonable amount of time. It also means less waiting around in line for scarce supercomputer time.

Marvin C. Wunderlich of the National Security Agency in Ft. Meade, Md., has used the MPP for factoring large numbers. His implemen-

tation adds 10 digits to the length of a number that can be factored by computer using the "continued fraction" method (SN:3/30/85,p.202). With this method, the MPP takes up to 14 hours to factor numbers carrying as many as 64 decimal digits. Other factoring methods do the job faster, but Wunderlich's work provides a benchmark for the speed that can be attained with a given number of parallel processors.

"The MPP has brought new life to the continued fraction method," says Wunderlich. In a computer with more processors than the MPP, this method may match the performance of the "quadratic sieve" algorithm, which has been used to crack the largest numbers to date.


Other scientists have used the MPP to build ocean circulation models, to study how charged particles interact with the earth's magnetic field, to draw topographic maps from satellite photographs, to model groundwater movement, to reconstruct X-ray images and for many other scientific applications.

In general, MPP users have found two serious problems. First, the "host" VAX computer is too slow for efficiently getting data into and out of the MPP. One researcher calculates that although the MPP itself can do 6.5 billion additions in 1 second, it would take half an hour to read in that many numbers. The second problem is that individual processors have too little local memory, which limits the complexity of operations that can take place at each processor.

The MPP can be upgraded fairly easily, says Goodyear's Ken Batcher, who designed the machine. Various faster devices for controlling the flow of data can be used to bypass the VAX computer, he says, and memory chips that hold more data can be substituted for those now in place. However, Batcher wryly notes, "no matter how big you build the memory, some program will always overflow it."

"Some things will change," says James R. Fischer, head of Goddard's image analysis facility, "and some things are much too expensive to change." The MPP will probably get a new, faster host computer early this year, he says. However, increas-

The MPP has more than 2,000 of these custom-made very-large-scale integrated-circuit chips. Each chip contains eight processors.



Goodyear Aerospace

ing the amount of memory at each processor would cost nearly \$2 million.

Many users also point out the importance of immediate access to graphic images of computed results. "We need ways to see the results of our processing," says Goddard's James P. Strong. That allows researchers to try things out until they get the kind of result they want. It also helps them understand complicated sets of data. Several projects are now focusing on how the MPP can be used to produce graphic images directly and on ways to convert computed data efficiently into suitable images at remote graphics terminals. How quickly that develops depends on NASA's budget.

Researchers with new projects can still apply for time on the MPP, says Fischer. Moreover, almost all of the present members of the MPP working group plan to continue their projects. "We've had very few people in the working group drop out," he says.

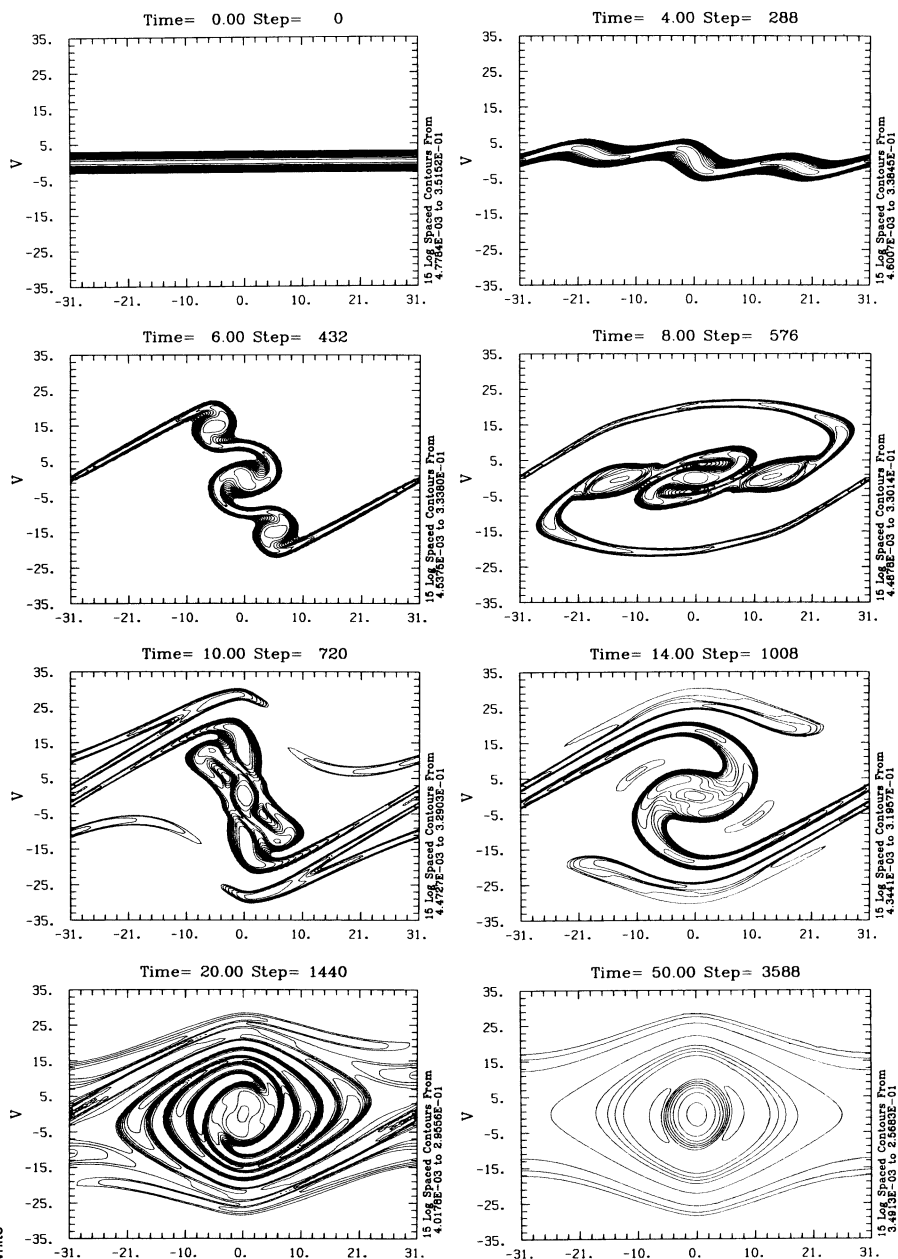
The MPP is not the only multi-processor machine now available. In the 1970s, International Computers Ltd. in England developed the Distributed Array Processor (DAP), which has 4,096 simple processors connected in a 64-by-64 array. As in the MPP, the processors all respond to the same instruction at the same time while working on different pieces of data. Unlike the MPP, however, the DAP provides electronic "highways" for sending data across the entire array instead of just to neighbors.

A new London-based company, Active Memory Technology, now proposes to build and market a cheaper, more compact, slightly simpler version of the DAP. When available next year, this 32-by-32 "personal parallel processor" will probably cost about \$150,000, says Dennis Parkinson, head of the DAP Support Unit at the University of London.

The Connection Machine, developed by Thinking Machines Corp. in Cambridge, Mass., and introduced as a commercial product last year after much experimental work (SN: 6/16/84, p.378), uses 65,536 linked processors to achieve high computing speeds. As in the MPP and DAP, all of these small processors run the same program simultaneously. Unlike its cousins, the Connection Machine's pattern of links between processors can be changed to fit the problem it is working on.

All three machines work best when relatively simple operations need to be performed on huge amounts of data. This comes up most often in applications like image processing. The Connection Machine has also been used for analyzing seismic data and for calculating fluid flow around objects.

The Connection Machine, DAP and MPP belong to a single family of parallel processing computers — those



The gravitational collapse of a one-dimensional system consisting of infinite sheets of stars, as it would appear with velocity plotted against position.

with a large number of simple processors connected in a relatively straightforward way. But there are others, and more are coming. Each machine has somewhat different capabilities and characteristics, strengths and weaknesses.

For example, the new "T Series" computers from Floating Point Systems, Inc., in Beaverton, Ore., also have a "massively parallel" architecture, but they compute with as many as 16,384 linked "transputers," sophisticated, high-speed microprocessors manufactured by a British company. Recently, the Encore Computer Corp. of Marlborough, Mass., received a three-year, \$10.7 million contract from the Defense Advanced Research Projects Agency to develop a new, massively parallel computer system that can execute 1

billion instructions per second.

Many other parallel processing architectures have been suggested and developed. Some of these machines use fewer but much more sophisticated and powerful processors, sometimes linked in exotic ways (SN:2/16/85,p.104; 2/23/85, p.117). Computer memory can also be allocated differently.

Which parallel processing computer to use depends on a blend of cost and efficiency. Different architectures suit different problems. The MPP, as one symposium participant remarked, was an attempt "to see how far you can get with lots of cheap, dumb processors." Despite its limitations, the MPP has shown that for some applications, it can run with the best of them. □