

SCIENTIFIC AMERICAN PRESENTS

# EXTREME ENGINEERING

**EXCLUSIVE**  
Photos from the  
World's Largest Dam

*see page 14*

DISPLAY UNTIL FEBRUARY 23, 2000

[www.sciam.com](http://www.sciam.com)

## MEGAPROJECTS

- the **TALLEST** building
- the **FASTEST** computer
- the **LONGEST** tunnel
- the **STRONGEST** bridge

Flying at  
**MACH 20**


Atomic  
Tweezers

Designing  
New Life-  
Forms

QUARTERLY \$6.95









# BLITZING BITS

by W. Wayt Gibbs

**S**troll into your local computer store, plunk down \$2,000, and you can take home a fairly zippy machine. With a seventh-generation 600-megahertz processor, 128 megabytes of memory and about 13 gigabytes of disk capacity, a late-model personal computer can tear through that full-screen, full-motion, post-apocalyptic shoot-'em-up with scarcely a hiccup. But snazzy computer games are one thing. Simulating what actually happens inside a detonating nuclear bomb—or a collapsing star or a folding protein—requires a qualitatively different kind of machine, a machine that has not yet been built. It took the Blue Pacific system at Lawrence Livermore National Laboratory,

currently the fastest supercomputer in the world, 173 hours to complete the turbulence simulation shown here. Your state-of-the-art PC would have to hum along for well over 16 years to do the same job, assuming it worked at its peak speed of 600 megaflops (million floating-point operations per second)—which, of course, computers never do. Every 20 days you would have to add another 13-gigabyte hard drive to store the results.

And yet, says Mark K. Seager, one of the supercomputer gurus at Livermore, this massive computation will shed light on just one small, idealized part of the problem. To confidently answer whether refurbished bombs will burst, how the globe will warm, how the universe took shape, and other questions that elude theory and experimentation, scientists need computers of 10 to 10,000 times the speed and capacity of Blue Pacific. One such machine is already under construction, and others are on the drawing board.

**VIRTUAL VISCOSITY:** In the biggest simulation of its kind ever attempted, the Blue Pacific supercomputer calculated how two adjacent fluids would mix after one was hit by a shock wave moving at Mach 3. Two thirds of the supercomputer's 5,760 IBM processors worked on the problem, performing 308 quadrillion calculations. The results—a series of 27,000 three-dimensional images (one portion of one thin slice of one image is shown here)—filled 3.76 terabytes of disk space.




# Inside the Fastest Computer

**B**lue Pacific, the U.S. Department of Energy's 3.9-teraflops (trillion floating-point operations per second) supercomputer, became fully operational in May 1999. Just over one year later it will be made obsolete by the next step in the DOE's Accelerated Strategic Computing Initiative (ASCI). In a giant lab in Poughkeepsie, N.Y., IBM engineers are constructing a \$100-million, 10-teraflops successor, called ASCI White, which will occupy a large room (*below*) and


is scheduled for demonstration in March 2000 and for operation by late summer. Like Blue Pacific, the new machine will divide up programs to run on thousands of processors simultaneously.

ASCI White is but the fourth of seven supercomputers planned by the initiative, which aims to produce a 100-teraflops system by 2004. That level of performance, says Livermore's Mark Seager, is "the absolute minimum" needed to simulate how an entire nuclear weapon would detonate—or not.

**1** In January 1997 the sun expelled a wave of plasma that collided with Earth's magnetosphere, creating stunning auroras and providing a rich set of scientific observations that may help explain how our planet's magnetic shield works. To compare the observations with current theory, physicists at the University of Maryland created a detailed supercomputer simulation of the event, computing a kind of three-dimensional movie from the basic laws of physics.



**2** On ASCI White, such a simulation would first carve out an imaginary block of space around Earth. In order to perform the calculations as quickly as possible, the software would divide this volume (*above*) into perhaps 10 billion smaller "cubes." These units, each containing only a few mathematical operations and some initial numbers, would stream out of random-access memory (RAM) and move through a switch.






**3** The heart of any supercomputer is the switch (*eight dark-blue boxes shown above*) that pulses data between and among processors, memory chips and disks. ASCI White uses a "multistar omega network," which connects the 8,192 processors in the machine with one another and with 10,752 external disk drives in such a way that any processor is never more than two hops away from any other one. The switch can move data to and from each group of processors at a rate of 800 megabytes per second—more than five times the speed of Blue Pacific's switch.




# of Tomorrow



**6** After a week or two of around-the-clock operation, the final movie of 50,000 frames—a total of perhaps 500 trillion cubes—will be complete (*right*). To store such massive amounts of data, ASCI White will boast 195 terabytes of external disk storage (*left*). By way of comparison, the printed contents of the Library of Congress comprise about 10 terabytes.



**5** As the cubes of information enter the nodes, they flow into memory and are distributed among the processors. If all the numbers are in place, the processor can do its mathematical work, filling the cube (*left*) with the results and sending it back out over the switch to be stored in the disk farm. Often, however, the data in a cube are in RAM or are spread among two or more processors that must pass messages to one another to cooperate in arriving at an answer. This process slows the computation enormously, and as a result supercomputers rarely operate at more than 20 percent of their theoretical peak speed.



**4** The processors are organized into nodes, which are grouped four to a case. Every node (*left*) in turn houses 16 375-megahertz POWER3-II microprocessors. This chip is designed to execute four floating-point calculations simultaneously, for a peak performance more than twice that of a 600-megahertz Pentium III. In addition to more than eight megabytes of cache RAM per processor, every node contains at least eight gigabytes of local memory and two internal 18-gigabyte hard drives.

DAVID FIERSTEIN



# Three Ways to Get from Here

**F**or any other field of engineering, the idea that one could achieve a thousandfold increase in performance in less than a decade would be sheer lunacy. We will not soon see cars streaking along the desert at 700,000 miles (1.1 million kilometers) per hour or buildings that rise 280 miles into outer space. But many computer scientists believe that by 2007 they will be able to build a supercomputer that delivers one petaflops (a quadrillion floating-point operations per second)—three orders of magnitude faster than Blue Pacific.

If it were made in the mold of ASCI White, a petaflops machine would contain at least 250,000 microprocessors, draw one billion watts—the output of a large nuclear power plant—and cost roughly as much as a fleet of aircraft carriers, estimates Thomas L. Sterling of the California Institute of Technology. Its

processors would waste most of their time waiting for data to arrive from memory. To skirt these obstacles, researchers are investigating at least three radically different designs.

**1** A consortium led by Caltech is working on the most ambitious of the three approaches to petaflops capability, a so-called hybrid-technology multithreading architecture. “Hybrid technology” means that the researchers intend to use new kinds of chips, networks, disks—everything. Massive microchips will have logic circuits woven among large banks of memory, so that the two can communicate more rapidly. The “single quantum flux” chips will be cooled to near absolute zero so that they superconduct and use about one millionth the energy of conventional processors. This also should allow them to run at speeds exceeding 150 gigahertz, so that “only”

2,048 processors are needed for the system’s multithreaded operation, in which a program is broken into individual tasks that can be performed concurrently.


Information will flow through the system as light in optical fibers rather than as electricity in copper cables, increasing bandwidth by a factor of 100 or more. And up to one petabyte (million billion bytes) of data will be stored as holograms in crystals rather than as magnetic patterns on spinning disks, greatly boosting speed and reducing power consumption.

The consortium has prototypes of some of the components, but a number of the technologies are still in the research stage. Nevertheless, with backing from four federal agencies, this effort is the best funded of the petaflops supercomputer designs.

**2** The first computer to break the teraflops barrier was no bigger than a large photocopy machine, and it sat in a humble room at the University of Tokyo. Called GRAPE-4, it did only one thing—calculate the gravitational attractions among many objects such as stars or asteroids—but it performed its job with exceptional efficiency, surpassing on that narrow range of problems the speed of even the mighty Blue Pacific.

GRAPE-6, now nearing completion, should set another milestone, hitting 200 teraflops by executing sizable chunks of its program on special-purpose microprocessors. This is a cheap way to build supercomputers if you need them for just one kind of problem, says Mark Snir, manager of scalable parallel systems at IBM. “We looked here at what it would take to build a multipetaflops machine customized to the problem of protein folding,” he says. “We could do it for a few million dollars—much, much less than a general-purpose machine.”

There may be a way to have the best of both worlds. Recent generations of so-called configurable chips—processors



COURTESY OF DAVID PORTER, University of Minnesota

**STELLAR SECRETS:** The seething interior of a young star (modeled here by researchers at the University of Minnesota who used a Silicon Graphics Origin2000 supercomputer with 128 processors) is one of many mysteries that petaflops-speed supercomputers may one day unravel.




# to Petaflops

that can rewire their circuitry on the fly—raise the hope of supercomputers that can transform themselves into ultrafast machines custom-designed for the problem at hand. But configurable chips are still so slow and expensive that the idea remains little more than a hope.

**3**As of late August, the fastest and hardest-working computer system on the planet was not behind razor wire at a classified lab or humming in the bowels of some university building. It cost less than \$1 million to set up and almost nothing to run. It never had to come down for maintenance, and it grew faster every day.

SETI@home, a small program written at the University of California at Berkeley and distributed over the Internet, was released this past May. Within three months, more than a million people had downloaded the software, which scans signals recorded by the Arecibo radio telescope in Puerto Rico for signs of extraterrestrial intelligence. With SETI@home installed, each PC downloads from Berkeley a chunk of data to process, performs the calculations while the



EICHIRO KOKUBO University of Tokyo; HITOSHI MIURA Musashino Art University

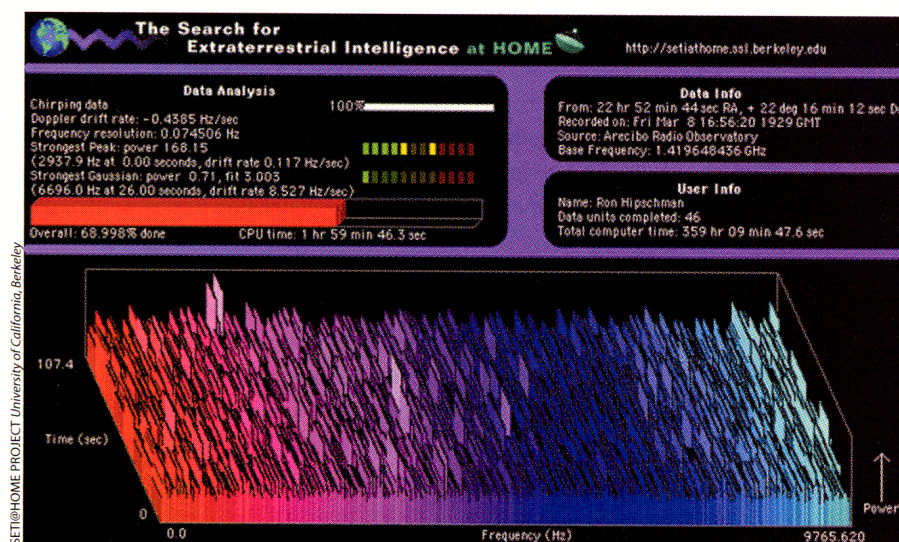
**LUNAR BIRTH:** The origin of the moon (*gray cluster at top left*) was simulated in record time and detail by the HARP supercomputer, a predecessor of the GRAPE-4 machine that used special-purpose chips to calculate the gravitational attractions among many objects, including stars and asteroids.

machine would otherwise be idle and then sends the results back.

By September the results were pouring in at the rate of seven teraflops. Put another way, a popular screensaver had in four months zipped through computations that would have taken the Blue Pacific supercomputer about 26. This

may be a special case, points out Dan Werthimer, the project's chief scientist. "I don't think we could attract one million people in 224 countries to help with one of the 'grand challenge' problems," such as turbulent mixing.

But at Berkeley and elsewhere, SETI@home does provide inspiration to researchers who are trying to build "virtual supercomputers" by connecting, say, all the computers in a university or a hospital and harnessing processing power that would otherwise go to waste. That research raises the possibility that one day in the near future, the Internet will offer a way not merely to communicate but also to tap into a nearly unlimited reservoir of computing power. SA



**VIRTUAL SUPERCOMPUTER:** The screensaver SETI@home has enabled more than one million PCs to join the search for extraterrestrial life.

## About the Author

W. WAYT GIBBS is senior writer at *Scientific American*. He has completed some 30 units for the SETI@home project but has discovered no E.T.s phoning home.