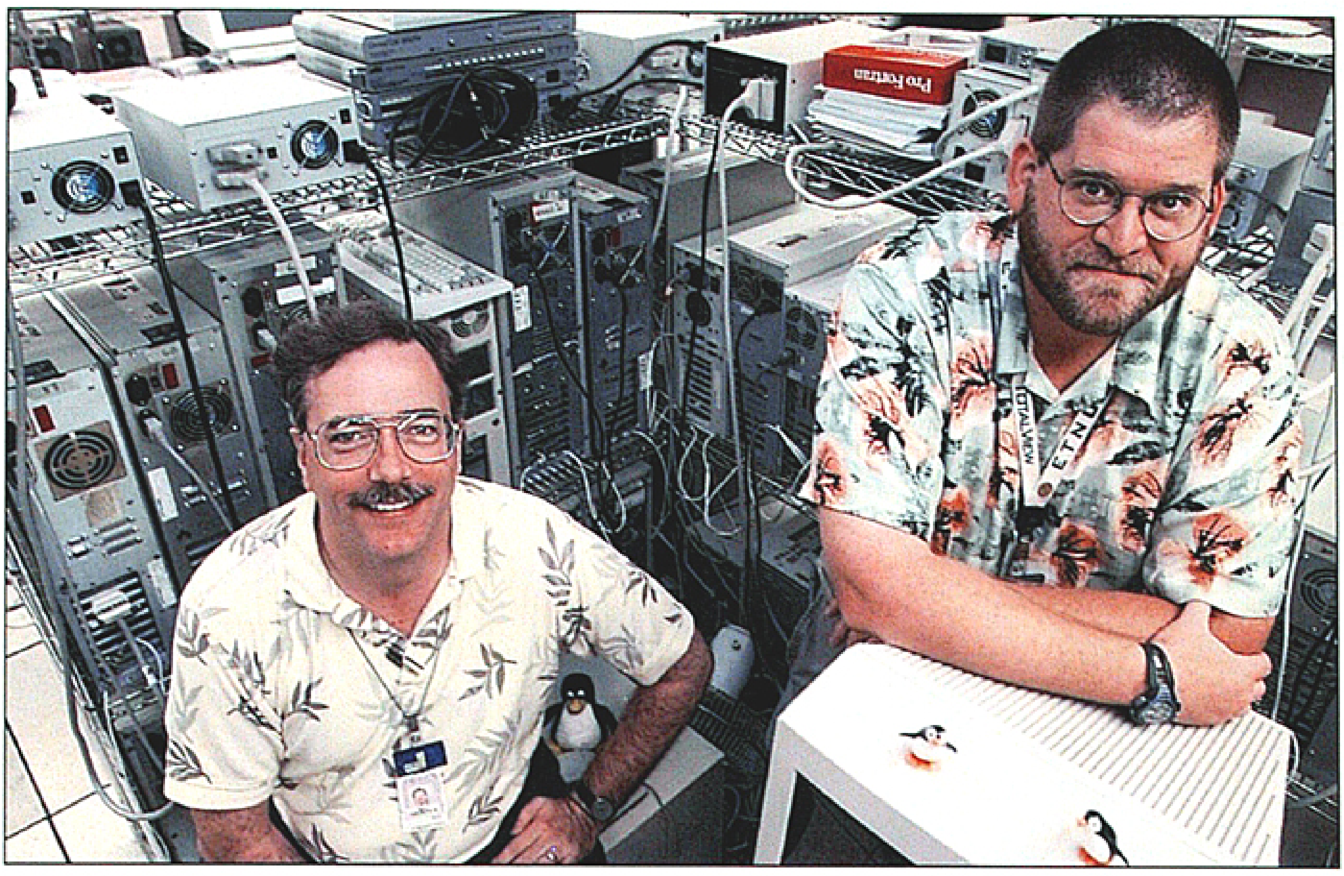

Bill Hargrove, left, and Forrest Hoffman use computer clusters to conduct environmental research at ORNL. The scientists write computer codes that enable the old PCs to do parallel processing tasks similar to supercomputers.

Analytical horsepower

Two scientists at Oak Ridge among pioneers in harnessing humble PCs into data dynamos

Cluster computing is like harnessing a team of horses to pull a heavy load.

That's an analogy offered by Bill Hargrove and Forrest Hoffman, who were among the first to effectively turn groups of PCs into research workhorses.

Before assembling their first cluster of computers in 1996, the Oak Ridge scientists were drowning in ecological data. They were desperate for a quick way to analyze many sources and types of information.

"Now we just want more data," Hoffman said during a visit to their makeshift computer room on the first floor of ORNL's environmental sciences building. Computers were stacked everywhere, in racks and on the floor.

Their initial venture was dubbed the Stone SouperComputer, taking its name from the fable about a wanderer who promises to feed a village by boiling water and adding a stone. Village skeptics ultimately turn to volunteers, contributing vegetables and enough fixings to make a hearty soup, thus proving that good things come from meager beginnings through cooperation.

At ORNL, Hargrove and Hoffman built their cluster by mixing and matching dozens of castoff PCs from the lab's "swap shop" and gifts from cooperative secretaries who upgraded from 486s to better machines.

With virtually no capital investment, the Stone SouperComputer grew to 128 nodes (single processors) and was able to accomplish parallel-processing tasks equivalent to a small supercomputer.

"It really comes down to your codes," said Hoffman, a computer specialist in the environmental sciences division. "We invested a lot of time writing programs for dynamic load-balancing. That assigns more work to the faster nodes and less work to the slower nodes, but together they all combine to solve the problem."

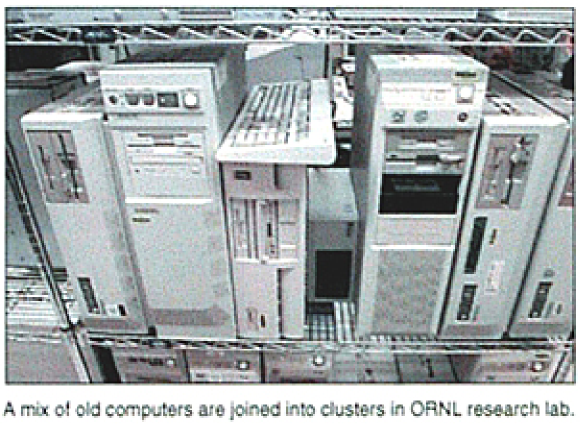

A mix of old computers are joined into clusters in ORNL research lab.

"WE'VE BEEN ABLE TO TACKLE BIGGER PROBLEMS WITH HIGHER RESOLUTION THAN PEOPLE WOULD EVEN THINK OF DOING BEFORE."

— Bill Hargrove, landscape ecologist

The parallel PCs helped the ORNL team produce detailed maps of U.S. "eco-regions," with areas color-coded according to their environmental similarities. The computer nodes evaluated and correlated ecological data with 25 or more variables, everything from nitrogen levels in the soil to the humidity in April.

Hargrove, a landscape ecologist, said, "We've been able to tackle bigger problems with higher resolution than people would even think of doing before."

The Stone SouperComputer is now partially disassembled, and Hoffman and Hargrove no longer rely on junk computers to do their work. They have a budget to buy new processors needed for projects that support the laboratory's work on climate change.

The cluster they use now, also known as a Beowulf, named for the one developed at NASA, cost about $50,000 to build. It has nine nodes, each of which has two Intel Pentium III processors running at 1.0 Ghz; 1 GB of memory (RAM), 18 GB of disk space, and two Fast Ethernet interfaces for file-sharing and log-in access.

"While this system is relatively small, its capacity exceeds that of the Stone SouperComputer when we shut it down," Hoffman said.

"The great thing about these Beowulf-style clusters is that we can add to them when we need more capacity and upgrade them slowly over time as projects and problems necessitate."

A science fad of the 1990s is now a reliable, everyday tool at research institutions big and small.

That's due, in part, to software development at ORNL that makes it relatively simple for scientists everywhere to hook up a bunch of PCs and operate them as a single parallel computer.

"What makes it so popular is the cost," said Al Geist of ORNL's Center for Computational Sciences.

Instead of investing millions of dollars in a high-performance supercomputer, a university can build a powerful research cluster for $150,000.

Scientifically useful groupings can be had for a lot less. The cost of setting up a cluster is about $1,000 per node.

Even the smallest of colleges are using them to teach students how to do parallel computing and to do their academic research, Geist said.

The Oak Ridge lab heads a national consortium that developed OSCAR (Open Source Cluster Application Resources), the most popular cluster-management software program. According to Geist, it's "brain-dead easy" to use. A CD downloads the software that shows how to connect the machines and make them work together.

"We're very proud of the fact that we're making it possible for scientists all over the world to group PCs into clusters for good use," he said.

Though not officially classified as supercomputers, big clusters can be formidable research machines.

In order to do research on software and other cluster developments, ORNL's computing center has assembled its own 64-processor system known as Extreme TORC. When it's not engaged in cluster development, the experimental system is available to staff researchers.

Cluster computing is highly individualized. Hoffman and Hargrove develop their own research tools, implementing software from ORNL or elsewhere. They also use computer clusters to prepare some projects to be run on IBM supercomputers available at the lab.

"If the problem is small enough, we'll do it here on the cluster because we don't have to wait in line," Hoffman said. "If it's a really enormous problem — if we can't run it in two days or less — then we'll do the prep work here and migrate the data over to the IBM, do the runs there, bring the data back and do the hard-core analysis here."

Hoffman writes a monthly column for Linux magazine called "Extreme Linux," in which he discusses how to build and run these machines and how to write programs to maximize their capabilities.

Geist said cluster computers are now used in almost all fields, from oil exploration to airplane design.

"The biggest one in the world is probably run by Google, which has about 6,000 processors," he said. "But they don't do any interesting science. They're just using it to regurgitate URLs."

It's not unheard of for a research institution to have a cluster sitting right next to its supercomputer, Geist said. But, even at their best, cluster systems can't interact as well as a high-end supercomputer specifically designed for many processors to work together, he said.

And, while clusters can boost a scientist's research capabilities, connecting four or more PCs at home would probably offer few benefits to Average Joe or Average Jane.

"Most folks who use their computer to surf the Web and check their e-mail already have a computer that's probably 10 times more powerful than they need," Geist said.

Even video games, the single biggest power draw on most home computers, wouldn't benefit much from a cluster because those games typically are set up to run on a single processor, he said.

Lab linking up to universities, other centers

Oak Ridge National Laboratory wants to expand its network capacity, establishing computer links with other research institutions around the nation.

Much of this will be accomplished through cooperative efforts with ORNL's university partners, particularly the core universities – Duke, Georgia Tech, Virginia, Florida State, Virginia Tech and North Carolina State – that were part of UT-Battelle's original proposal to manage government facilities in Oak Ridge.

"We are the nation's steward in some sense, so we always have to be thinking of how to make these machines available to researchers across the nation," said Thomas Zacharia, the laboratory's associate director for computational sciences.

In mid-August, the laboratory announced that it had linked its supercomputers to the Atlanta Giga-Pop, the high-speed access point for Southern universities.

The data pipeline is hundreds of times faster than the fastest home dial-up connection.

ORNL Director Bill Madia, in keeping with the Atlanta theme, said the new computer link could download the DVD version of "Gone With the Wind" (all 222 minutes of it) in 6 seconds.

Charles Liotta, vice president of research at Georgia Tech, who manages the Atlanta Giga-Pop, wryly added, "It took only 2 seconds to do the latest Austin Powers movie."

Liotta predicted that high-speed computer connections to Oak Ridge resources would have big impacts not only on research, but education and economic development as well.

"The outcome will allow the Southeast to prosper," he said.

Len Peters, the vice provost for research at Virginia Tech, said efforts are under way to provide a similar pipeline from ORNL to the Mid-Atlantic Crossroads in Washington, D.C. That would make the Oak Ridge resources available to many other universities.

"We need these networking roots to get us down to Oak Ridge with high speed," Peters said. "Access is very important."

Ray Orbach, director of the U.S. Department of Energy's Office of Science and former chancellor of the University of California-Riverside, said economic development is based on scientific discovery, and advanced computing will play a critical role in the 21st century.

Orbach said people need to realize that "teraflops" and "petaflops" are more than just fancy lingo.

"If we are going to make scientific discovery using simulation as our vehicle, we need to be operating in this new environment," he said during a visit to ORNL.

Just as important as high-speed computing is the ability to get data back and forth between the research institutions, the DOE official said. That enables the sharing of top research tools, such as the Spallation Neutron Source under construction in Oak Ridge, he said.

The SNS, scheduled for completion in 2006, is designed for experiments to be performed remotely.

"There is not any difference between a control room in Oak Ridge and a control room in Atlanta or anywhere else in the world," Orbach said. "We will literally be able to do those experiments in real time. ... But the only way we do that is if the data is going back and forth sufficiently fast."

Future

Continued from R1

he's proud he played a role in building the program's reputation. But the UT official said the turnaround at ORNL occurred nearly a decade earlier, in 1985, when then-Director Herman Postma decided to establish a supercomputing capability.

Instead of trying to compete with other labs for the best technologies that existed at the time, Postma chose to focus on the future and jump a generation ahead. As a result, ORNL was able to acquire Serial No. 1 of Intel's first commercial parallel computer, an IPSC/860.

"It was in the same ballpark as a supercomputer, and it was powerful enough to do some serious work," Ward said.

Years later, Oak Ridge continues to pursue an individual strategy in supercomputing.

Unlike other top computing centers, ORNL limits the uses and users for its top-of-the-line supercomputers, and thus gets more bang on its research projects — especially the task of modeling the world's changing climate.

Lawrence Berkeley National Laboratory in California, for instance, accommodates about 2,000 users a year at its computer research center. Oak Ridge is equipped with similar capabilities but limits its user list to about 200.

Zacharia said the top machines are reserved primarily for research in biology, climate change and advanced materials, and that approach won't change.

"If anything, we're going to be even more focused on tackling really important challenges," he said. "We don't want to be like a honeybee going from one flower to another. We want to stay the course until we have an impact and then go on to another field."

ORNL has had an association with most of the major computer makers.

The laboratory, for instance, recently has been working with IBM on a project known as Blue Gene, as well as having two IBM machines in the supercomputing center.

Zacharia notes that Oak Ridge worked years ago with Cray, then acquired a big machine from Intel and then IBM and then Compaq and now is preparing to receive the latest from Cray once again.

"Our task," he said, "is not to be the private-sector company building machines or selling machines. Our task is doing science, and assembling the best computers to do the science. That's my job."

ORNL is not interested in exclusive arrangements with any company, he said.

Zacharia predicts the lab will have a petascale computing capability — 1,000 trillion calculations per second — before the end of the decade. That's hundreds of times greater than what's now available.

"We will get there. If not the leader of the pack, we will be in the lead pack," ORNL's computing chief said. "I have to pinch myself about what I just said. It's not that far away. It keeps me awake at night."

ORNL through the decades: THE FIFTIES

1950 Oak Ridge School of Reactor Technology (ORSORT) established, and two new test reactors begin operation.

1951

1952

1953 Engineers design transportable reactor for the Army to use in remote sites such as Antarctica.

ORACLE, world's most powerful computer, starts up.

1954 Engineers working on nuclear-powered aircraft use the Tower Shielding Reactor to support design of shielding to protect a nuclear airplane's flight crew from radiation.

1955 Alvin Weinberg named laboratory director, a position he would hold for 18 years.

1956 First experimental bone-marrow transplants performed in mice.

1957

1958

1959 Sen. John F. Kennedy and his wife, Jacqueline, visit laboratory.