Campus 'supercomputer' helps decipher mysteries of science

by David F. Salisbury

Figuring out how circuits in the brain work, identifying why some drugs are more effective with certain people, determining the properties of a state of matter that existed milliseconds after the creation of the universe, modeling the folding and unfolding of proteins and how that affects the way they work, and designing electronic devices that can survive the high radiation of outer space. These are a few examples of the projects researchers across campus are pursuing with the aid of Vanderbilt's home-brew computer cluster – the total computing capacity of which makes it one of the 500 most powerful supercomputers in the world.

This spring, the facility processed its 3 millionth job, racking up a total compute time of 14.5 million hours – the equivalent to a single PC grinding away for more than 1,600 years. Not bad for a system that began in 1994, when Vanderbilt physicists cobbled together 50 workstations to create a primitive network that allowed individual computers to work in parallel.

In the late 1990s, Professor of Physics Paul Sheldon and Professor of Molecular Physiology and Biophysics Jason Moore (who has since gone to Dartmouth) joined forces to spearhead construction of a true parallel-processing machine. VAMPIRE, as the Vanderbilt Multi-Processor Integrated Research Engine was dubbed, was a 55-node cluster built in April 2000 from off-the-shelf parts by a group of 20 physicists, biologists and computer technicians in a marathon, two-and-a-half day stretch.

The project not only served as a magnet for computer-power-hungry researchers from a variety of disciplines, but also fostered a number of cross-disciplinary collaborations as researchers discovered that computational procedures developed in other disciplines can help them solve their own problems.

Within a few years, demand for VAMPIRE's services outstripped its capacity. In 2003, two researchers from the School of Engineering – Professor of Electrical Engineering Ron Schrimpf and Peter Cummings, the John R. Hall Professor of Chemical Engineering – rode to the rescue by obtaining funding to add a second, 120-node cluster. Later that year, Vanderbilt's administration invested $8.3 million to make the facility into a university-wide resource for researchers from all disciplines and, in the process, gave it the less colorful moniker of the Advanced Computing Center for Research and Education (ACCRE). Combined with a $1.5 million equipment grant from the National Institutes of Health obtained under the leadership of Moore and Professor of Molecular Physiology and Biophysics David Piston, the cluster was expanded to its current size of 1,500 processors.

Deciphering brain circuitry

Jeffrey Schall, the E. Bronson Ingram Professor of Neuroscience, is one of the psychology department's heavy computer users. He studies the circuitry in the brain that controls eye movements. For example, in 2000 he published the discovery of the "oops center" in the brain, an area that is activated when the brain realizes it is making a mistake. He has measured activity levels in thousands of individual neurons in this area in the brain of the macaque and is using the data to figure out exactly how the brain processes information.

“By probing the patterns of activity in the circuits of the brain, we can get new information about how the brain works,” Schall said. “We know the brain is something like a computer, but it is a ‘black box’ that we don’t understand. So we put it through its paces and measure its performance. At the same time, we probe the circuit elements inside to see if we can reverse-engineer it.”

One of the major obstacles Schall is trying to overcome is “model mimicry,” the fact that there are a number of plausible ways that the brain can process information. “We couldn’t do the simulations required to address this issue without the cluster,” he said.

Probing gene-gene interactions

Marylyn Ritchie is an assistant professor of molecular physiology and biophysics who received her doctorate from Vanderbilt in 2004. Her research area is computational and statistical genetics, and an important element in her decision to stay at Vanderbilt was access to ACCRE. She is studying the way in which genes interact with each other and with environmental factors to cause disease and affect therapies.

Human DNA contains about 30,000 genes. There is general agreement that the most common human diseases are not caused by single mutations but by variations in several different genes. However, the most current statistical methods can only identify with confidence genes that have strong individual effects. In addition, most illnesses are affected by environmental factors, which further complicates matters. As a result, there is a great need for new methods of analysis.

That is exactly what Ritchie is working on. She is collaborating on a large project called Platform for the Analysis, Translation and Organization of Large-Scale Data, or PLATO for short. It is an attempt to reduce the statistical complexity of the problem of identifying multiple-gene interactions by integrating existing biological information into the process. In her research, Ritchie is concentrating on environmental and drug effects, “because these are much easier to change than genes.”

Characterizing quark-gluon plasma

Physics Professor Charles Maguire has been using the ACCRE cluster to analyze an exotic new state of matter, called the quark-gluon plasma, which may exist briefly in the heart of exploding stars. More importantly, the entire universe may have existed in this state shortly after its birth in the Big Bang.

As a member of the Vanderbilt team on PHENIX, the largest of four detectors on the $600 million Relativistic Heavy Ion Collider (RHIC) at Brookhaven National Laboratory, Maguire supervised a group of 40 scientists from around the world that developed the software used to analyze the data produced by the house-sized instrument. The collider accelerates gold nuclei to 99.99 percent the speed of light and then crashes them together. The resulting collisions create microscopic explosions that dissolve the protons and neutrons, leaving a cloud of fundamental particles called quarks and gluons. These then recombine into hundreds of different short-lived sub-atomic particles that fly outward into the detector. In 2005, the researchers announced that the RHIC had not only created the quark-gluon plasma, but that this state of matter behaved like a perfect liquid instead of like the perfect gas they had anticipated.

In previous experiments such as RHIC, massive atom smashers would run for a brief period and generate such a large amount of data that it would take the physicists a year or more to analyze it. This made it very difficult to fine-tune the collider operation.

Today, RHIC researchers transport the data via the Internet to ACCRE and other similar computer clusters and analyze the data in near-real time. “It’s a tremendous advantage,” Maguire said. “It allows the collaboration to monitor the different physics signals on a daily basis and then fine-tune the detector to enhance the science content during the remainder of the accelerator run.”

Unwrapping protein origami

In the sub-cellular world of proteins, form follows function. Proteins are made of different combinations of 20 standard units called amino acids. However, a protein’s biological function is determined more by its physical shape than its chemical structure. Figuring out the various ways that proteins fold and unfold in different conditions is a critical step in figuring out how they work and interact.

“There are a lot of proteins whose amino-acid sequence has been identified but whose physical structure is unknown,” said Jens Meiler, assistant professor of chemistry and pharmacology. “That is particularly true of the membrane proteins that are the targets for more than half of all drugs.”

Creating accurate, three-dimensional models of membrane proteins is one of the problems that Meiler is combating with the supercomputer power of ACCRE. He is a member of a major collaboration that is developing an algorithm, called Rosetta, designed to accurately predict protein folding. Using Rosetta, he has successfully built a model of a crucial protein, called a transporter, that binds with serotonin and allows neurons and other cells to store the neurotransmitter involved in such diverse functions as learning, sleep and the control of mood. The transporter itself is a target for cocaine and other illicit drugs. Working with the laboratory of Randy Blakely, the Allan D. Bass Professor of Pharmacology, Meiler is validating the model by predicting how tightly different molecules will bind with the transporter and then checking the predictions experimentally.

“With ACCRE, we can do the model calculations much quicker than the people in the lab can run the experiments to test them. That is definitely an argument in favor of the cluster,” Meiler said.

Outer space device design

Vanderbilt’s Institute for Space and Defense Electronics is widely acknowledged as the leading center for research on the effects of radiation on electronic devices. Its researchers are heavy ACCRE users.

“If you use electronics in a space system, one of the biggest issues is that they will be exposed to radiation because they are no longer protected by Earth’s atmosphere,” said Ron Schrimpf, the professor of electrical engineering and computer science who directs the institute. “There are two general categories of problems that occur. One is the long-term degradation of parts due to the accumulation of the energy that radiation deposits on them. The other is loss of data caused by the deposition of charges large enough to upset the bits.”

Currently, the susceptibility of electronic devices is determined by testing them in particle accelerators, a lengthy and expensive process. Schrimpf and his colleagues are developing software tools that can replace this procedure with “virtual testing.” They have developed simulations of the radiation environment that run on the supercomputer cluster. “We can specify that a system is operating in geosynchronous orbit or that there is a major solar flare – whatever consideration we might want to take into account,” Schrimpf said. They also simulate how specific devices respond when radiation creates electrical charges in different locations.

Among other things, this allows the researchers to test the extent to which changes in various circuit parameters – such as resistances, capacitances and spacing between individual elements – can harden a circuit against radiation. “Formerly, people looked at a handful of different variations and guessed what might be best. Now, with ACCRE, we can run 25,000 different variations overnight and have the answers back the next day,” Schrimpf said.

posted 04/30/07