

Organizations


U.S. Department of Energy
Oak Ridge National Laboratory
Oak Ridge, TN
www.ornl.gov



Purdue University
Lafayette, IN
www.purdue.edu

Cray XT5™ “Jaguar” supercomputer

With a peak speed of 2.33 petaflops, “Jaguar” is a Cray XT5 supercomputer located at Oak Ridge Leadership Computing Facility (OLCF). With the ascension to the petascale, Jaguar is able to give computational scientists unprecedented resolution for studying a whole range of natural phenomena, from climate change to energy assurance, to nanotechnology and nuclear energy.



Schematic view of a nanowire transistor with an atomistic resolution of the semiconductor channel.




Illustration of electron-phonon scattering in nanowire transistor.

ORNL and Purdue Explore Technology at the Nanoscale

A team led by Gerhard Klimeck of Purdue University has broken the petascale barrier while addressing a relatively old problem in the young field of computer chip design.

Using Oak Ridge National Laboratory’s Jaguar supercomputer, Klimeck and Purdue colleague Mathieu Luisier reached more than a thousand trillion calculations a second (1 petaflop) modeling the journey of electrons as they travel through electronic devices at the smallest possible scale. Klimeck, leader of Purdue’s Nanoelectronic Modeling Group, and Luisier, a member of the university’s research faculty, used more than 220,000 of Jaguar’s 224,000 processing cores to reach 1.03 petaflops.

The team is working to help manufacturers transcend Moore’s Law, which since 1965 has anticipated the astonishing pace of technology advances. It was in that year that Intel cofounder Gordon Moore suggested in Electronics magazine that the density of transistors on a computer chip — and, therefore, the speed of that chip — would double about every two years. And indeed they have, with Moore’s company just this year releasing chips that hold more than 2 billion transistors on a piece of silicon just over an inch square.

An unavoidable reality

But technology makers are running across an unavoidable physical reality: There is, in fact, a limit to how many transistors you can fit onto a sliver of silicon. Within the next few years transistors will be as small as they can get. “We’re at the stage where these transistors — the on-off switches — in these chips are as small as 20 to 30 nanometers in some of their critical widths,” noted Klimeck. “That’s along the lines of 100 to 150 atoms. We’re beginning to reach a limit where making it another factor of two smaller is going to be more and more difficult.”


Twenty nanometers is indeed very small. By comparison the shortest wavelength of visible light — belonging to violet — is 20 times wider at 400 nanometers. An especially fine strand of hair is 1,000 times wider at 20 micrometers.

The power of Jaguar has given Klimeck and Luisier the ability to pursue this task with unprecedented realism. “What we do is build models that try to represent how electrons move through transistor structures,” Klimeck explained. “Can we come up with geometries on materials or on combinations of materials — or physical effects at the nanometer scale — that might be different than on a traditional device, and can we use them to make a transistor that is less power hungry or doesn’t generate as much heat or runs faster?”

“You’re reaching the stage where you can’t think of your material being like a continuum. We experience that this material is not continuous, but discrete. And the placement of these atoms is important.”

NEMO and OMEN

The team is pursuing this work on Jaguar with two applications, known as Nanoelectric Modeling (NEMO) 3D and OMEN (a more recent effort whose name is an anagram of NEMO). NEMO 3D is an evolution of the earlier NEMO 1D, which was developed at Texas Instruments in the mid-1990s to model resonant tunneling diodes — devices that used the quantum mechanical ability of electrons to tunnel through a tiny barrier and appear on the other side. NEMO 3D expanded the



System Overview

- ➔ Peak Petaflops: 2.33
- ➔ Six-core AMD Opteron™ processors: 37,376
- ➔ AMD Opteron cores: 224,256
- ➔ Compute nodes: 18,688
- ➔ Memory (TB): 300
- ➔ Dish bandwidth (GB/s): 240
- ➔ Disk space (TB): 10,000
- ➔ Floor space (SqFt): 5,000

“Having machines like Jaguar made these calculations possible.”

—Gerhard Klimeck
Nanoelectronic Modeling Group
Purdue University

Corporate Headquarters
Cray Inc.
901 Fifth Avenue, Suite 1000
Seattle, WA 98164
Phone: 206-701-2000
Fax: 206-701-2500
www.cray.com

application to three dimensions and million-atom systems, but at a cost: The systems were static, with no electron flow.

“We were able to say where the electrons would be sitting in this geometry, but we couldn’t afford computationally to model electrons injected at one end and pulled out at the other end, which we did do in NEMO 1D but in a one-dimensional space or representation,” said Klimeck. “Around 2004 compute powers became strong enough on massively parallel machines that we could dream of combining the NEMO 1D and NEMO 3D capabilities and actually represent the device in three dimensions one atom at a time and actually inject electrons at one end and pull them out at the other end. Having machines like Jaguar made these calculations possible. The theory of how to do that was understood with NEMO 1D, but it was computationally prohibitively expensive. OMEN is the next-generation prototype that runs on Jaguar now.”

These applications combine techniques to get the greatest possible scientific discovery out of modern supercomputers. They calculate the most important particles in the system — valence electrons located on atoms’ outermost shells — from their fundamental properties. These are the electrons that flow in and out of the system. On the other hand, the applications approximate the behavior of less critical particles — the atomic nuclei and electrons on the inner shells. The applications solve the venerable Schrödinger equation, which describes the quantum mechanical properties of a system, although Klimeck’s team has modified the equation to allow for electrons moving into and out of the system. Using OMEN Luisier has been able to model about 140,000 atoms. Without taking into account the flow of electrons, NEMO 3D is typically used to simulate 5 million to 10 million atoms.

At this scale the quantum mechanical oddities of electrons — such as the tunneling behavior studied in NEMO 1D—become increasingly important. “Our understanding of electron flow in these structures is different,” Klimeck explained. “As you make things very small, you expose the quantum mechanical nature of the electrons. They can go around corners. They can tunnel. They can do all kinds of crazy stuff. So electrons are quantum mechanical particles at that scale, and the atoms are discrete.”

Bringing insights to the real world

The team is working with two experimental groups to bring this research to the real world. One is led by Jesus Del Alamo at the Massachusetts Institute of Technology, the other by Alan Seabaugh at Notre Dame.

With Del Alamo’s group the team is looking at making the electrons move through a semiconductor faster by building it from a material called indium arsenide rather than silicon. With Seabaugh’s group the modeling team is working on band-to-band-tunneling transistors. These transistors bear some promise in lower-voltage operation, which could dramatically reduce the energy consumption in traditional field-effect transistors.

While there may be limited opportunities to decrease the size of a transistor, it is possible to make devices smaller and more efficient. One problem that has not been overcome is the energy lost as heat, which leads to challenges in both powering devices and cooling them.

“Atoms like to vibrate in certain ways,” Klimeck explained, “and the vibrations are losses that come out of the device as heat. If we understand how the electrons interact with the crystal at the nanometer scale, how heat is being generated, maybe we can avoid the heat and make the transistors less power hungry. On the flip side we may be able to improve devices that convert thermal energy into electricity.”

As he notes, minimizing the power requirements and heat generation of electronics as well as harvesting energy are critically important as our lives become more and more automated.

“This is important anyplace where you don’t want to lug 40 pounds of batteries with you, for instance, if you’d like your iPhone to be lighter and run all kinds of image processing on it,” he said. “This is true for any mobile or computing technology — anything where you don’t want to waste energy, which these days is anywhere.”

Content courtesy of Oak Ridge National Laboratory. Research sponsored by the DOE Office of Science.