Chinook: EMSL's Powerful New Supercluster

There are three pillars of science—experiment, theory, and simulation. While there is no question that these pillars are intertwined in major scientific advances, it is not often that all three are found within a single organization or facility. Fortunately, national and international researchers can find all three at the Department of Energy's Environmental Molecular Sciences Laboratory (EMSL).

EMSL is one of about 45 user facilities in the United States (sidebar "EMSL: A National Scientific User Facility"). Researchers from universities and national laboratories pair the first two science pillars in EMSL's problem-solving environment by using one-of-a kind equipment to generate more—and, in some cases, better—scientific data than ever before. The third pillar, simulation, also is available at EMSL in the form of a new supercomputer, named Chinook (figure 1; sidebar "Chinook: Connected to the Northwest" p62).

Chinook is a Hewlett-Packard (HP) 163 teraflop/s supercluster based on AMD Opteron processors with an InfiniBand interconnection network. HP assembled it using largely commodity hardware and software. There is sufficient flexibility in this supercomputing approach to allow a system to efficiently and effectively run demanding EMSL science applications without resorting to expensive special-purpose hardware. This approach works at a variety of scales and is becoming increasingly popular.

EMSL: A National Scientific User Facility

The Environmental Molecular Sciences Laboratory (EMSL), is a DOE national user facility supported by the Office of Biological and Environmental Research and located at Pacific Northwest National Laboratory. Its purpose is to provide researchers worldwide with integrated experimental and computational resources for discovery and technological innovation in the environmental molecular sciences to support the needs of DOE and the nation. Since opening its doors in October 1997, EMSL has enabled the research of thousands of users from all 50 states and nearly 30 countries who represent academia, industry, and the national laboratory complex.

No scientific user facility can support high-impact science without investments in cutting-edge instrumentation and expertise to support visiting researchers. In addition to its $24 million Hewlett-Packard-developed supercomputer, Chinook, and other advanced computational resources, EMSL offers users access to state-of-the-art capabilities in:

- mass spectrometry
- nuclear magnetic resonance spectroscopy
- microscopy
- kinetics and reactions
- deposition and microfabrication
- spectroscopy and diffraction
- subsurface flow and transport

Users are encouraged to combine one or more of these capabilities to solve research challenges critical to the nation in four rapidly growing areas: biological interactions and dynamics; geochemistry/biogeochemistry and subsurface science; science of interfacial phenomena; and the emerging area of atmospheric aerosol chemistry.

Researchers can access the instruments and expertise through a peer-reviewed proposal process, which includes proposal calls under its science themes and computationally intensive research; access is usually free if the work will be published in open literature. Researchers also may submit a general-use proposal to EMSL at any time.

Further Information
http://www.emsl.pnl.gov

PNNL

Figure 1. *The Department of Energy's EMSL supercomputer, known as Chinook, is helping users of a national scientific user facility advance molecular science in areas such as aerosol formation, bioremediation, catalysis, climate change, hydrogen storage, and subsurface science.*

When designing a computer architecture for chemistry and biochemistry simulations, which account for about 90% of Chinook's work, it is essential the computer system has the right balance with respect to the processor, memory hierarchy, interprocessor communication, and disk access and storage. For most high-performance computing systems it does not make sense to invest in fast local storage, as few application codes outside of chemistry benefit from it. On Chinook, however, fast local storage is essential to create the best hardware balance that minimizes the time-to-solution for researchers running computational chemistry codes. To achieve the best balance for chemistry on Chinook, HP added an unprecedented amount of local storage bandwidth that temporarily stores data while the system works on other calculations. This resource allows for faster processing of complex codes, which is ideal for researchers desiring fast time-to-solution for their chemistry-related simulations.

Research projects in chemistry, biology, and environment are already realizing the benefits of Chinook's unique architecture. One team has used Chinook for large-scale calculations of the movement rate of molecules. They got their results within days, instead of weeks or months, which is how long it took previous generations of supercomputers. The team believes the project could not have been run even a few years ago, because the processing power for such large chemistry-based equations did not exist.

Chinook also is well-suited for non-chemistry work, such as analyzing vast amounts of biological genomics data. While researchers may not utilize all of Chinook's capabilities when running non-chemistry problems, those running genomics analyses certainly benefit from its fast interconnect and shared storage. In addition, because Chinook's architecture consists of mostly commodity components, it is relatively easy for researchers to transfer previously-developed cluster codes to the supercomputer.

Because Chinook's architecture consists of mostly commodity components, it is relatively easy for researchers to transfer previously-developed cluster codes to the supercomputer.

Chinook: Connected to the Northwest

When EMSL awarded the contract for developing its next-generation supercomputer to HP in September 2007, staff and users alike felt tremendous anticipation in obtaining a machine that would help researchers address complex environmental problems quicker and with larger and more realistic models.

When staff at the Pacific Northwest National Laboratory in Richland, Washington—where EMSL resides—were asked what name EMSL's new supercomputer should be given, they responded with a landslide of suggestions, many related to the environment.

The name "Chinook" emerged as the clear winner. Chinook is a type of salmon that swims in the Pacific Ocean, as well as the rivers of the northwest, with strong ties to EMSL's location in the Pacific Northwest.

Chinook is also a Pacific Northwest Indian word meaning "snow eater," and the term is used in the region to reference the warm, wet—and most welcomed by some people—coastal winds that, when strong, can make one foot of snow almost vanish in a day.

So, whether it be fish or wind, EMSL has reeled in (or is it blown in?) a new system that will help advance scientists' research in areas such as aerosol formation, bioremediation, catalysis, climate change, hydrogen storage, and subsurface science.

Chinook is No Small Fry

Chinook's architecture (figure 2; sidebar "The Guts of Chinook") has unique components that are necessary for chemistry applications, yet the system is mainstream enough to be accessible to a variety of users, across a broad spectrum of disciplines.

Chemistry calculations have relatively small-sized data input, but they need large—sometimes huge—intermediate storage while they are running on the computers. With speed, memory, and storage requirements in mind, HP added architectural features to Chinook that are inimitable within the supercomputing community. For example, the feature that most sets it apart from other supercomputers is its 800 terabytes of local scratch disk space. It is increasingly rare for compute nodes to have local disks at all, and Chinook has eight in each of its 2,300 nodes.

The extra storage provides an alternative to physical memory; 18,480 disk drives (sidebar "Reliably Running Thousands of Disk Drives" p64) make it possible for the supercomputer to pre-compute intermediate terms within a calculation and then store them on a disk, rather than re-computing them every time they are needed—something that would cost a lot of extra compute cycles and time. As Chinook continues to run the calculations, it pulls the pre-computed integrals from the disk. The result is a more efficient run of calculations and less time-to-solution. Local scratch largely remains the most cost effective way to provide the shear amount of bandwidth that is required to sustain efficient large-scale computational chemistry simulations. Another critical element of Chinook's hardware is its aggregate local disk bandwidth of 924 gigabytes per second. Put another way, all projects that run on Chinook have more than 400 megabytes per second of dedicated disk bandwidth on each compute node. This amount is in contrast to other supercomputers, which have little to no local storage at all.

Each node within Chinook has two quad-core AMD Opteron processors, 16 gigabytes of RAM, 350 gigabytes of local disk space, plus InfiniBand Host Channel Adapter, which enables swift communication between nodes. The combination of these attributes is what makes Chinook a perfect fit in an environment such as EMSL.

Computational chemistry problems require the ability to move big datasets across nodes in short amounts of time, which Chinook can do with its InfiniBand interconnect with Double Data Rate, 4x InfiniBand connections. Each connection is a set of parallel communication channels or "lanes," and each lane provides five gigabits per second of bandwidth. So, Chinook's four-lane (4x) connection provides 20 gigabits of bandwidth. By design, an InfiniBand network implements a "fabric" of connections between compute nodes. Every compute node can use one of many paths through the fabric to reach any other node. This feature can be used to prevent network congestion by routing traffic to other nodes, or to allow nodes to still reach each other if part of the network fails. By comparison, Ethernet allows nodes to use only one path between any two nodes, so congestion can be more of a problem and network failures can easily isolate nodes from each other.

The InfiniBand fabric (figure 3) is made of silicon chips that can make 24 connections to other chips or hosts. Thirty-six of these chips make up

Chinook's computing power is roughly the same as 9,200 desktop computers working together in your home or office.

The Guts of Chinook

Chinook was designed with unique components that are necessary for chemistry applications, yet the system is versatile and accessible to a variety of users across a broad spectrum of disciplines.

- 163 teraflop/s of peak performance
- 2,310 computing nodes
- 4,620 Quad-core AMD Barcelona 2.2 gigahertz processors
- 36.96 terabytes of memory
- InfiniBand DDR 4x Interconnect
- 1.3 petabytes of scratch disk space
- 924 gigabytes per second of local disk bandwidth
- 270 terabytes of global file system
- 40 gigabits of network connectivity

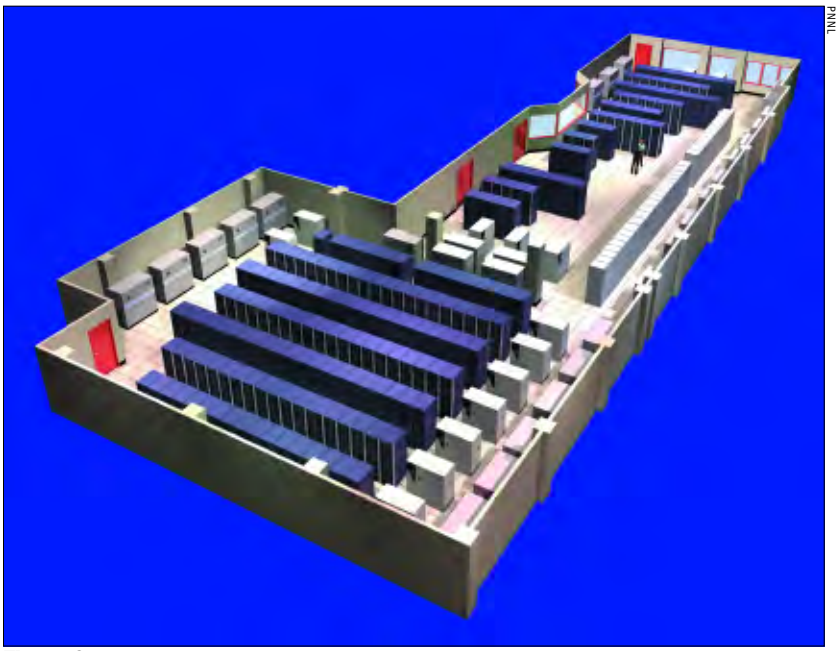

PNNL

Figure 2. *A three-dimensional rendition of the Chinook supercomputer at EMSL.*

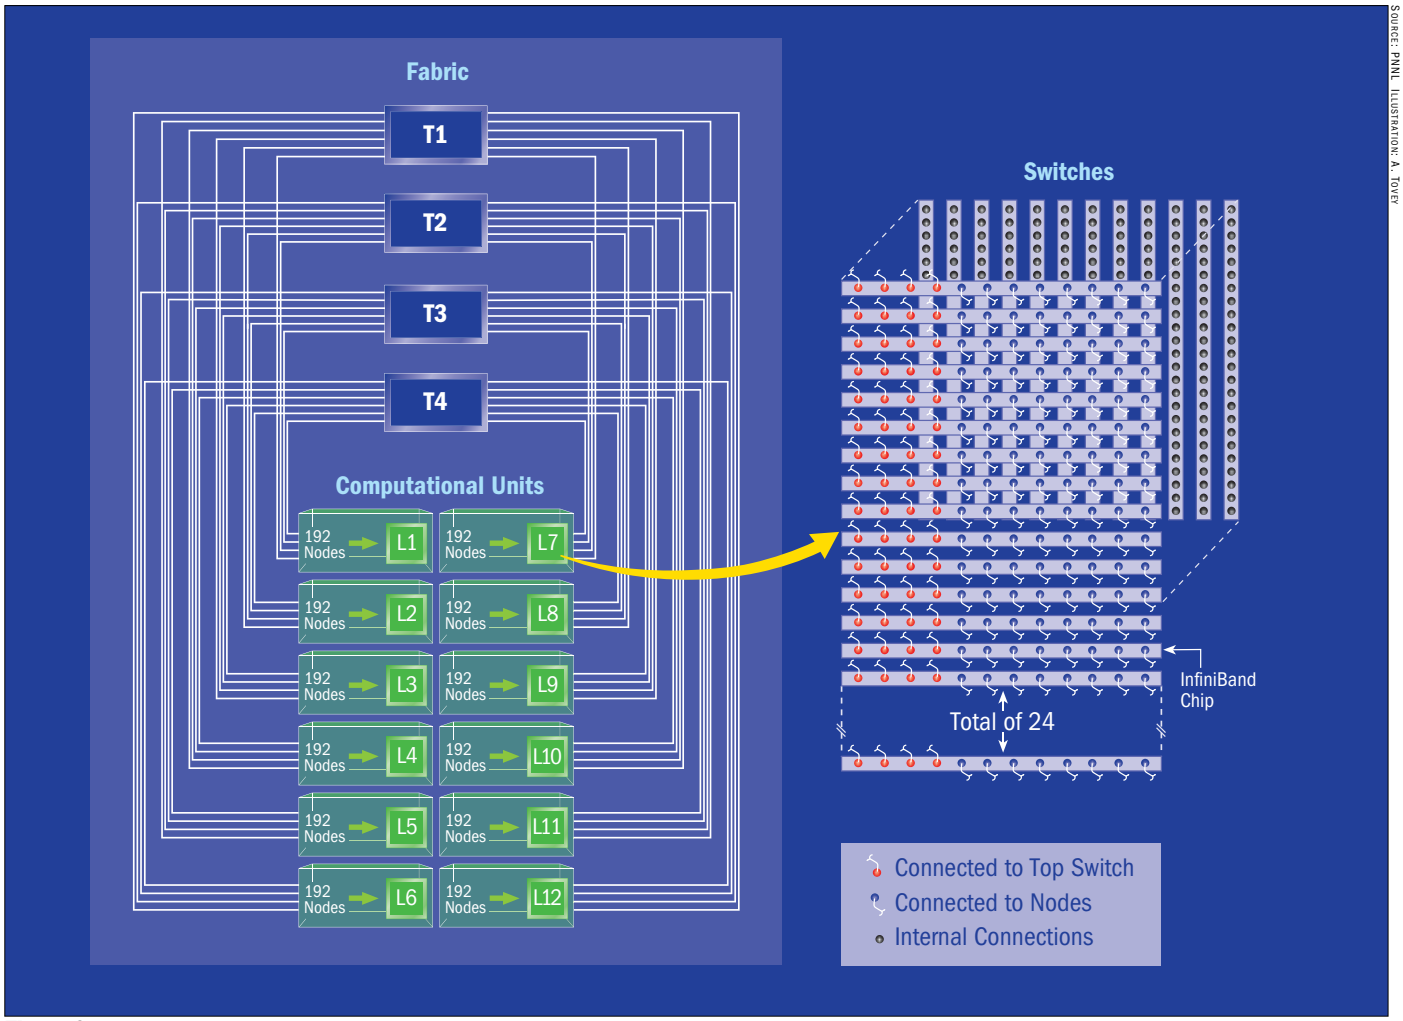

SOURCE: PNNL ILLUSTRATION: A. TOVEY

Figure 3. *On the left, arrangement of the four top (T) switches and 12 Leaf (L) (one per CU) switches that make up Chinook's InfiniBand fabric. Each of the lines represents 24 physical connections. On the right, cable arrangement showing the 24 InfiniBand chips in each switch. The arrangement with eight ports connected to nodes and four ports connected to top switches allow messages to cross the fabric with five hops or less. The 12 chips (depicted vertically) are embedded in the switch and allow full bandwidth connectivity between nodes in a given CU.*

Reliably Running Thousands of Disk Drives

All storage devices in Chinook employ some degree of redundancy to ensure that no single disk failure interrupts jobs or the system. Compute nodes have two hardware RAID (redundant array of inexpensive disks) controllers with four disks each. Data are striped across a set of four disks in a RAID5 configuration, which allows the storage to continue to function in the event of a disk failure. The two sets of RAID disks are glued together using software. When a disk fails in a compute node, the system allows whatever current job is running on that node to finish, then marks it for repair. Once a hardware technician replaces the drive, restoring the node to full functionality, it is marked eligible to take on new work.

one of Chinook's 288-port InfiniBand switches. Twelve of these switches have 192 or 194 compute nodes connected to them, and the remaining 94 or 96 ports are connected to four top-level switches. The top-level switches allow communications to be routed up and down the fabric so that every node has fast communications with every other node. For each chip that communications must pass through, a delay—or latency—of about 150 nanoseconds is added. One of many goals within high-performance parallel computing is to minimize latency, so it is important to minimize the number of "hops" through chips (sidebar "NWPerf").

To minimize hops on Chinook, HP organized its compute nodes into computational units (CUs) that are built around a single large InfiniBand switch. As a result, nodes within the CU take no more than three hops to communicate with another node in the same CU. In fact, sometimes it may take only one hop to go from node to node. Cluster Resources, Inc. designed Chinook's scheduling software, Moab, in such a way that the software knows about the system's CUs and has features to schedule jobs within a single CU if the user prefers to minimize latency in that way. If a job cannot be scheduled within a single CU, the schedule can spread it across multiple CUs. In the latter case, though, communications between any two nodes may take three or five hops.

Many research projects that run on Chinook will also benefit from its global scratch that uses HP's Scalable File Share (SFS) filesystem—Lustre under the covers. This scratch is particularly handy for computations that generate large datasets or require comparisons between large datasets because they need high-speed, shared storage in order to run efficiently. With 270 terabytes of scratch and the ability to sustain 30 gigabytes per second of read/write activity, Chinook can efficiently and effectively handle such calculations. SFS has come full circle since it was tested and refined on Chinook's predecessor at EMSL before HP began selling it commercially.

Nearly 23,000 disk drives provide a terabyte per second of disk bandwidth—fast enough to transfer each second a volume of data equivalent to 126 DVD movies, which would take nine and a half days to watch.

Software is Key

Having the hardware is important, but having scientific applications that can fully utilize this hardware is just as vital. In computational chemistry, one of these highly-scalable scientific codes is NWChem.

NWChem is part of the Molecular Science Software Suite, MS3, which includes NWChem, the Global Arrays Toolkit, and the Extensible Computational Chemistry Environment (ECCE). Each has a distinct role, but when used together, the suite provides a comprehensive environment enabling scientists to perform complex chemistry calculations on a supercomputer.

NWChem is popular because it scales to thousands and, in certain cases, tens of thousands of processors. "NWChem is recognized as DOE's premier computational chemistry software and it continues to be used to solve complex scientific problems of importance to DOE," said Steven Ashby, Deputy Director for Science and Technology at Pacific Northwest National Laboratory (PNNL).

EMSL's high-performance software team has implemented an extensive array of cutting-edge computational chemistry methods in NWChem. It is intended to run on high-performance parallel supercomputers to address scientific questions that are relevant to dynamic and reactive chemical processes occurring in our everyday environment—for instance, catalysis, photosynthesis, hydrogen storage, protein functions, and environmental remediation. Understanding these processes will help researchers develop new materials for hydrogen storage, improve the efficiency of solar cells, and develop technologies to clean up contaminated soil.

Thousands of people worldwide use the software to investigate chemical processes on molecular systems that range in size from tens to millions of atoms. The software allows them to apply various theoretical techniques, ranging from highly-correlated methods to molecular dynamics, to predict the structure, properties,

NWPerf

EMSL has relied on a node-level performance characterization tool known as NWPerf (figure 4) to understand performance of scientific codes both at a per-job level and at a system-wide level. This understanding has enabled scientists to identify and fix performance anomalies of an individual code, and it provided the basis for setting the architectural balance requirements specified for vendors.

NWPerf was originally developed for MPP2, and then was re-engineered for use on Chinook. Part of the re-engineering was using the open-source Ganglia distributed monitoring system to collect data in place of the previous generation's custom data collection mechanism. Today, EMSL administrators and scientific consultants depend on NWPerf daily for real-time feedback on Chinook's performance.

EMSL is targeting summer 2009 for an Open Source release of NWPerf

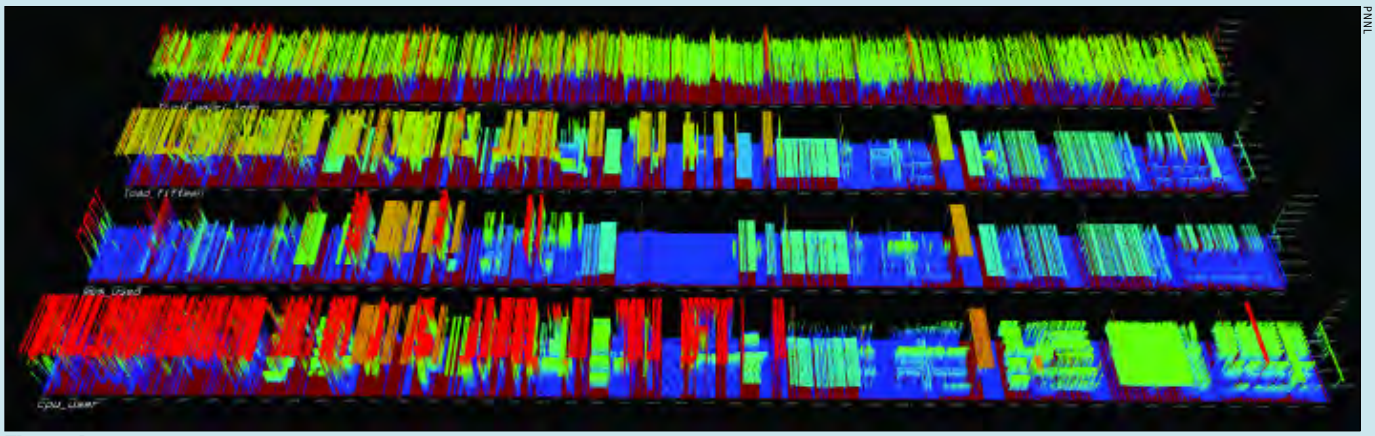

Figure 4. *NWPerf provides a graphic representation of Chinook's performance. Four graphs on this chart represent (bottom to top): cpu_user—time spent on the CPU running user code; mem_used—memory allocated by processes on the system; load_fifteen—average number of jobs in the running state over the last fifteen minutes; power_ambient—the temperature of the node in Celsius taken inside the node near the power supply. Each graph has one little line for each node in the cluster. So for cpu_user, each line shows the cpu load over the last 128 minutes, with "now" at the back of the graph, and 128 minutes in the past at the front of the graph. The color and height of the line show the usage, which in this case ranges from 0% (blue) to 100% (red) of the node. The red lines in the graph show nodes that are busy running user code.*

and reactivity of chemical and biological species. The software is free to users and the team provides expert support. While it runs efficiently on the large supercomputers in the DOE complex, NWChem also can run well on conventional workstations or mid-range clusters.

The team of developers makes sure the software is scalable both in its ability to solve large problems, as well as its usage of available parallel computing resources. For example, on EMSL's previous supercomputer, researchers were able to run calculations at 63% of peak efficiency utilizing all 1,840 processors, while also utilizing the high-speed interconnect and memory available. At the National Energy Research Scientific Computing (NERSC) Center, the team has been able to utilize 8,096 processors to run *ab initio* molecular dynamics calculations, and they are pushing the software even further to use tens of thousands of processors at the same time.

For its parallel scalability, NWChem relies on the Global Arrays Toolkit—the second piece of MS3. This toolkit consists of high-performance, efficient, and portable computing libraries and tools that enable application software to run on a variety of parallel computing systems. The toolkit provides an efficient and portable shared-memory programming interface for distributed-memory computers using one-sided communication (where possible), a library for solving linear systems on parallel architectures (PEIGS), and a parallel I/O library called ChemIO.

The final element of MS3 is ECCE, which provides experimental and computational scientists with an easy-to-use, tightly integrated suite of graphical user interface and visualization applications for molecular modeling, code input generation, job management, and (real-time) output analysis. The data management framework in ECCE provides researchers with easy access to their computational data in an organized fashion. The software's simple graphical user interface makes it effortless for a broad range of scientists around the world to apply computational chemistry methods on supercomputers to their research (sidebar "The Petascale Data Storage Institute" p66).

"NWChem is recognized as DOE's premier computational chemistry software and it continues to be used to solve complex scientific problems of importance to DOE."

Steven Ashby
PNNL

The Petascale Data Storage Institute

The Petascale Data Storage Institute (PDSI), funded by DOE's SciDAC program, focuses on data storage problems found in petascale scientific computing. Three universities and five national laboratories comprise PDSI, which is led by Carnegie Mellon University.

Each organization brings a different expertise to the group, and PNNL's is its "storage story" and knowledge of Lustre.

With eight disk drives on each of Chinook's 2,300 compute nodes, the supercomputer has a scratch disk that is entirely unique in the U.S. supercomputing community. The PDSI team plans to run tests during system downtimes that push parallel file systems to performance levels and stripe widths well beyond what is possible on other systems.

Chinook's predecessor, MPP2, also had lots of local and scratch disks (about 7,000). The PDSI team carefully studied the records captured from those disks during their life. PDSI researchers discovered patterns of messages coming from drives 24–48 hours before failure. Similar data from Chinook will be captured and studied, hopefully increasing the ability to see into the future far enough to proactively repair hard drives before they cause an unplanned interruption to a simulation.

SOURCE: PNNL ILLUSTRATION: A. TOVEY

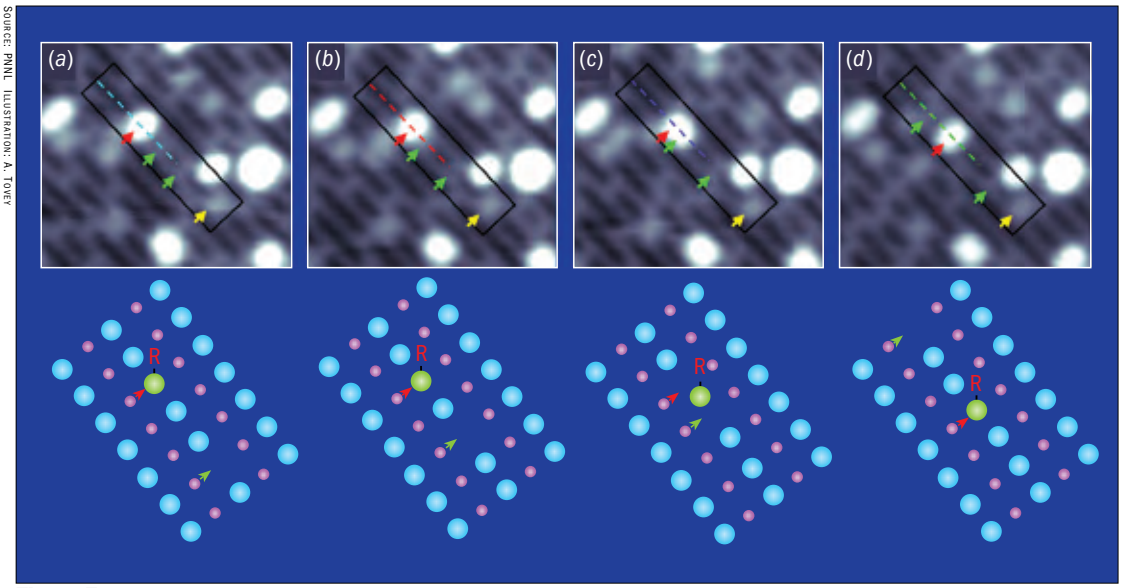

Figure 5. *Understanding how molecules, in this case alcohol, move on the surface of a titanium oxide catalyst. Images from a scanning tunneling microscope (top) and a simple model show how the oxygen from the alcohol (green with "R" on top) diffuse on the surface of titanium dioxide (the oxygen is represented by the blue spheres and the titanium by the purple). Massively parallel electronic structure calculations performed on Chinook were used to simulate these observations and provide a reaction mechanism compatible with the observed experimental kinetics.*

Chinook can do 163 trillion calculations per second. If you could sit down with pencil and paper and do one multiplication problem per second, every second, it would take over five million years to do what Chinook can do in one second.

Developing New Sources of Clean Energy

Unique scientific instrumentation, theory, and simulation each played an important role in one team's quest to learn how atoms behave on the surface of a common catalyst like titanium oxide. Their discovery could help them tailor molecular delivery systems, leading to the development of new, clean energy sources, including the generation of hydrogen to fuel cars, and the design of technologies that use titanium dioxide, such as air and water purifiers.

The team of researchers from PNNL, the University of Texas–Austin, and Southern Illinois University tried to understand how molecules, in this case alcohol, move on the surface of a titanium oxide catalyst. They used a state-of-the-art scanning tunneling microscope to conduct experiments and Chinook to run calculations to test their theories (figure5). The team observed that the alcohol molecules did not behave as expected.

Alcohol molecules "jumped" among holes—where oxygen atoms should be—on the titanium oxide's surface. To understand why the molecules jumped, the research team used Chinook to calculate the rate at which they moved. The computer did this by taking simultaneous snapshots of the molecule's movement from point A to point B to point C, and so on, searching for the jumping path that required the least amount of energy. This provided them insight into why the alcohol molecule

was jumping, and how they could change its behavior by modifying the catalyst's surface.

Once molecules and their reactions are controlled, they can be put to good use. For example, researchers could design a catalytic surface littered with vacancies to coach the desired portions of an alcohol to where some of it can be converted into hydrogen fuel needed for alternative energy sources in the home or car.

The team's simulation approach allows researchers to postulate ideas of what they think they saw in the experiment and confirm or dismiss it with calculations, resulting in theories that researchers had not previously considered, but that could be explored further with an experiment. This computation-before-experiment would not have been possible before the emergence of 21st century computing capabilities.


PNNL

Figure 6. *EMSL provides a unique environment for users to research a variety of oxide mineral films and interfaces, nanoscale materials, electronic and catalysis materials, microfabrication and microanalytical separations, and sensing.*

Rapid Time-to-Solution for Environmental Research

On and below the ground, molecular interactions often involve living cells. In order to understand and exploit such interactions, scientists need both dynamic chemical calculations, as well as computer-based genome database search capabilities. These resources available at EMSL (figure 6), were used by a research team to conduct the genome searches to answer the team's fundamental questions about protein–mineral interactions. Going forward in the research, the team will tap Chinook to perform even more searches, but in much less time.

The research team from The Ohio State University, Virginia Tech, and PNNL used NWChem to simulate how proteins interacted with metal oxide mineral surfaces. These simulations are important because they help the team better understand what interactions hold a microbe to a mineral surface, and how the microbe affects geochemical processes such as bioremediation and biomineralization. Understanding how those processes work may enable researchers to find ways to improve them and, for example, enhance cleanup efforts by using more effective microbes.

The simulations showed that a small peptide made of a specific sequence of five amino acids takes on a shape that matches the metal oxide mineral surface. In fact, the amino acid sequence provides a much needed "grip" on to the mineral surface. The research team searched databases of currently available microbial genomes to identify other proteins with the same amino acid sequence to see if they, too, could bind to a metal oxide mineral surface. The search turned up matches from two microscopic organisms that are known to interact with minerals.

One of the matches researchers found was in a bacterium that is known to bind and harvest energy from iron oxide minerals. The researchers believe that the gripping protein helps the bacterium attach itself to the surface and orient its molecular machinery such that energy harvesting can occur. In addition to iron, the bacterium can also use uranium and technetium for its energy production. As a byproduct, it converts these radionuclides from a very mobile species to mineral forms that are immobile; therefore, they are much less likely to be transported to our water supply where they would pose a risk to humans and wildlife.

The second match was with microscopic algae—commonly referred to as a diatom—that makes, or biomineralizes, its own ornate silica shells. These shells (a metal oxide mineral) offer protection and give the algae a competitive advantage, such as added chemical resistivity or maximizing the amount of sunlight available to them. This protection has allowed them to thrive in the world's oceans and fresh waters where they remove the carbon dioxide from the atmosphere. The diatom does not have one but many copies of the gripping pieces, which suggests that these small amino acid pieces are important in the formation of the silica shells in these microbes, perhaps by controlling the size and shape of their shells. Understanding how nature controls the making of silica shells will enable researchers to improve technologies that depend on the synthesis of small minerals having a certain size and shape, such as the fabrication of tiny electronic and mechanical devices.

Chinook will enable these researchers to examine interactions between other (and larger) pro-

The team's simulation approach allows researchers to postulate ideas of what they think they saw in the experiment and confirm or dismiss it with calculations, resulting in theories that researchers had not previously considered, but that could be explored further with an experiment.

It took about 24 hours to run a sequence that a smaller computer would have taken weeks to complete. Now, with the advancements in Chinook, researchers expect the system to perform similar runs in just a few hours.

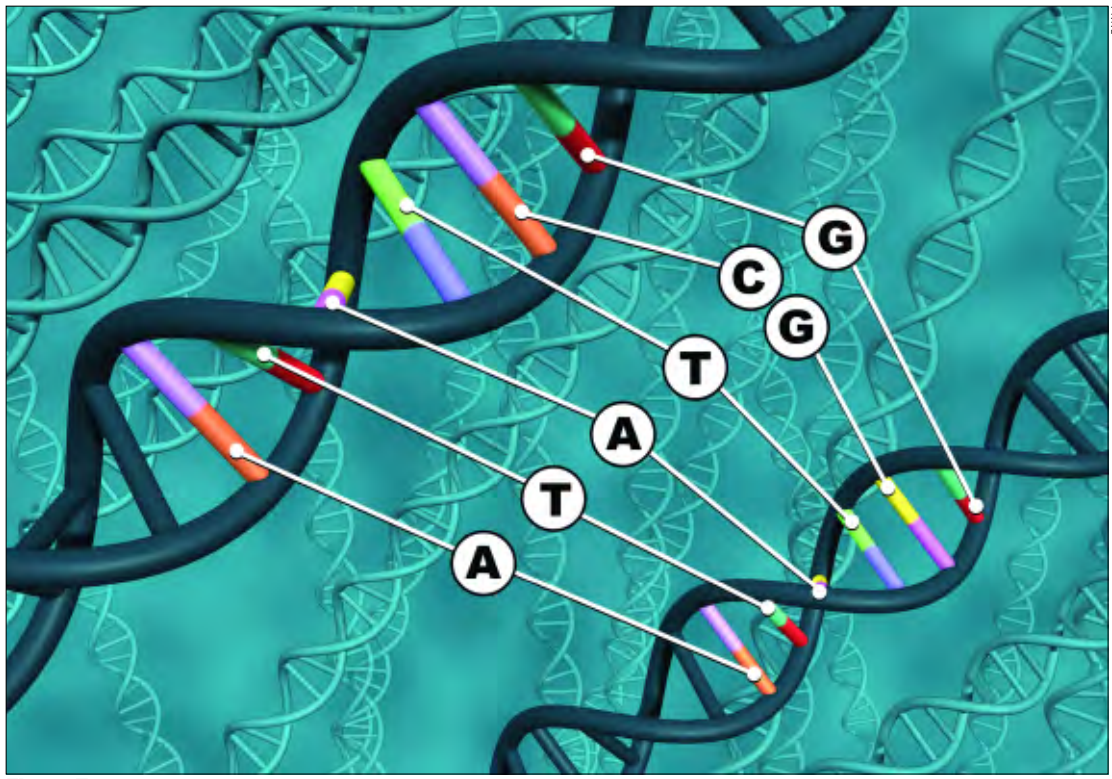

PNNL

Figure 7. *The goal of biosequence analysis is to identify DNA or protein segments that have similar chemical composition, and therefore are likely derived from a common ancestor. High-performance computing hardware and software is being used to drive this type of analysis at the scale needed for multiple genome analysis.*

teins and metal oxide mineral surfaces to determine which amino acid sequences bind, or bind better, to particular minerals.

A Big Picture of the Little Things

One of the advantages of integrating simulation and experiment is the ability to look at minute details on a very large scale. In computational biology, researchers analyze the sequence similarities of proteins so that they can make predictions about like proteins' functions.

Sequence analysis in computational biology can be extremely taxing on a supercomputer, especially when the analysis of the ever-growing data is nonlinear. For a supercomputer to successfully tackle such runs, it needs software that can help extract the biological information from the data, as well as a fast computer, which is accomplished on modern machines by using a large number of processors.

An example of such software is ScalaBLAST, which was developed at PNNL and is in use at DOE's Joint Genome Institute. It is a sophisticated sequence alignment tool that can divide the work of analyzing biological data into manageable fragments so that large datasets can run on many processors at the same time. The technology enables large computational problems to be solved in less than a day, rather than several weeks.

One such computational problem involves an all-to-all comparison of entire genomic databases—a nonlinear problem. With the comparisons, researchers are looking for how proteins change, or perhaps do not change, from one genome to another (figure 7).

In addition, the comparisons allow researchers to make predictions on the function of proteins based on the sequence similarity with proteins of known function. This knowledge helps researchers reduce the time and cost associated with learning about individual proteins. Moreover, knowing the functions of proteins can also provide researchers with insights into the function of the entire metabolic and regulatory pathways of which the proteins are a part.

Researchers used supercomputers throughout the DOE complex to run parts of this all-to-all comparison. It took about 24 hours to run a sequence that a smaller computer would have taken weeks to complete. Now, with the advancements in Chinook, researchers expect the system to perform similar runs in just a few hours. This increase in efficiency is largely due to the system's number of processors (going from 1,800 processors on MPP2, Chinook's predecessor at EMSL, to more than 18,000 processors on Chinook), the memory bandwidth, and the network bandwidth to move data quickly from node to node.

PNNL

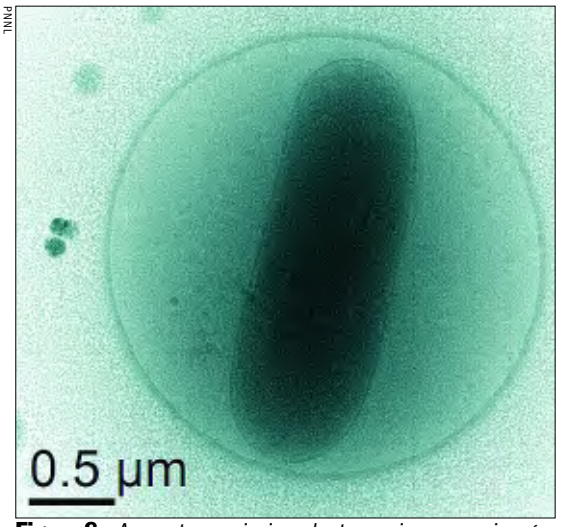

Figure 8. *A cryo-transmission electron microscope image of a bacterium with over 4,900 protein-coding genes stored in DNA molecules using 5,131,000 base pairs. Over 430 such genomes are in the database. If all 430 genomes had five million base pairs, that would require over 500 billion comparisons.*

Researchers in the field expect the computational needs of such bioinformatics problems to grow rapidly as new, more sophisticated instruments are producing sequence data at ever-increasing rates (figure 8).

Chinook Supports INCITE

Teams within the Innovative and Novel Computational Impact on Theory and Experiment (INCITE) program also benefit from Chinook's computing power. DOE's Office of Advanced Scientific Computing Research (ASCR), which sponsors INCITE, awards millions of supercomputer processor hours to aid researchers' work that involves analyzing, modeling, simulating, and predicting complex phenomena that are important to DOE.

Corning, Inc., a major producer of glass products ranging from liquid crystal displays and windows for space shuttles to optical fibers for telecommunication, leads one of the INCITE projects that is using Chinook. They are studying how organic molecules and oxide particles impact the flow of dense suspensions such as molten glass. Calculations to simulate these flow properties take large amounts of computer time as they try to elucidate how each type of organic additive interacts with each particle type. Their resulting impact on the flow properties of these suspensions in confined spaces has led to some surprises. What they are learning from the calculations gives them a good understanding of the physics at the atomic scale, which will allow them to better design future materials and improve manufacturing processes.

Looking to the Future

The state of supercomputing and its impact on science has changed dramatically over the past 20 years. The most notable change is the increasing integration of computation and experimentation to drive science. More and more, the scientific community is publishing research that relies on both experiments and simulations.

For example, most of today's high-end scientific instruments rely heavily on computational capability to interpret streaming data into a form and amount that researchers can use. These instruments are capable of generating terabytes of data per day, which a researcher with only a desktop computer could never analyze. Identifying the important bits from this sea of data, and using information gathered from multiple experimental approaches and simulation, requires high-performance computing.

Computation is evolving based on demands from the scientific community (sidebar "NWPerf" p65). As science demands faster systems with more storage, different architectures are being built, and more advanced software is developed to make efficient use of those architectures. Access to high-performance computing has become mainstream, rather than a novelty, for researchers. In part, this is due to advances in commodity cluster computing. As supercomputers are increasingly built from commodity technologies rather than proprietary technologies, there is more opportunity for sharing of resources.

The commodity cluster approach embodied by Chinook benefits the scientific community in multiple ways. The computational power is affordable enough that EMSL can offer its services to users without significant restrictions on the scale of jobs (that is, no maximum job size). In addition, Chinook is large enough that it can run project calculations to advance science, while continuing to help the computational chemistry community push its codes and science to a new level while preparing even larger simulations to be run by these codes on DOE's leadership class systems.

In the future, the relationship between experiment, theory, and simulation will only become stronger. National user facilities, such as EMSL, that already integrate these three pillars of science will continue to be valuable resources for users who are looking for these elements in a single place. The challenge for these facilities is keeping up with scientists' demands for sophisticated research equipment and robust computational systems that are necessary to advance the frontiers of science. Rather than just keeping up, though, EMSL aspires to stay ahead of the demands. •

In the future, the relationship between experiment, theory, and simulation will only become stronger.