



WAYNE STATE UNIVERSITY

High-Performance Computing Cluster Promotes Collaboration
Researchers line up to take advantage of powerful new parallel processing capabilities

High-Performance Computing Cluster Promotes Collaboration

Researchers line up to take advantage of powerful new parallel processing capabilities

By Art Mann, Silicon Mechanics

Researchers from various departments at Wayne State University (WSU) collaborate with one another frequently to work on innovative, interdisciplinary solutions to complex real-world problems. In order to efficiently test theories and models in complex systems such as cell injury in human tissue and defects in polycrystalline materials, these researchers need to utilize cutting-edge computing technology.

The university was recently awarded a high performance computing (HPC) cluster by Silicon Mechanics as part of its annual Research Cluster Grant program. The program, entering its fourth year, is designed to provide the latest in HPC technology to research institutions, thereby empowering groundbreaking research. The new cluster provides resources that are shared by a variety of the most computation-intensive research groups on campus, allowing researchers to break down boundaries of what research is possible.

A focus on bettering the community through research

The university had a substantial commitment to research prior to the receipt of this cluster, with annual research expenditures of nearly \$260 million. WSU is among only 3.5 percent of U.S. universities with the highest research classification from the Carnegie Foundation for the Advancement of Teaching. Interdisciplinary collaborative research teams, supported by WSU's Computing & Information Technology Department, include both computer scientists and domain scientists focusing on chemistry, mathematics, physics, biostatistics, and biology, along with cancer and biomedical researchers. They are often brought together by the need to incorporate parallel processing into their research to help advanced computational techniques perform better and faster.

Patrick Gossman, WSU's Deputy CIO, Community, Research and Special Projects, explained that a great deal of research and collaboration among the university's engineers, physicists, biologists, and medical doctors is moving towards "omics," the emerging

field of large-scale, data-rich collective characterization and quantification of pools of biological molecules that translate into the structure, function, and dynamics of an organism or organisms.

This is where the new HPC cluster comes in, as projects in this field require more computing capabilities than ever before. WSU already had a central HPC grid including more than 4,500 processing cores and 1.2 petabytes of data storage, but this grid does not offer the tremendous power of co-processors for highly parallel analyses. In addition, although a few researchers had a local server with one or two co-processors, there was no centralized infrastructure that combined the power of a cluster of processors with the power of multiple co-processors.

The new HPC cluster donated by Silicon Mechanics fills this gap, and combines multiple compute and GPU nodes to support the best of both types of parallel computing. The cluster is therefore a very significant addition to WSU's advanced computing environment. Gossman noted early benchmarking results showed a 90-fold increase in processing speed for one particular program on the new HPC cluster.

He added that using the same operating system, job scheduler and data storage for all advanced computing resources makes it easy for researchers to access the powerful new GPUs and Phi coprocessors.

The compute cluster was developed by Silicon Mechanics as a part of its annual Research Cluster Grant program, and includes hardware donated by Intel, NVIDIA, HGST, Mellanox Technologies, Supermicro, Seagate, and Kingston Technology, with cluster management software supplied by Bright Computing. The new cluster includes Intel® Xeon Phi™ coprocessors and NVIDIA® Tesla® GPUs.

University's technology leaders make a strong push for integrating advanced technologies with clinical and translational sciences

WSU's mission to be a premier urban research institution necessitates leading-edge computing technology

2014 Research Cluster Grant

so researchers can work toward solving the pressing economic and health problems that plague the world. Well known for both its engineering and physics departments, as well as being home to the nation's largest single-campus medical school, WSU has initiated a strong push to expand its related to integrative research teams tackling key health challenges incorporating basic discovery, translational, community engagement and implementation sciences focused on specific diseases and health care issues.

Such approaches can greatly benefit from computational and mathematical modeling of diseases and causative factors. Because of the complexity of the human body and the need to model and test clinical solutions before applying them to patient care, this field leans heavily on computational research and high performance computing.

Dr. Stephen Lanier, vice president for research, explained the rationale. "We don't want to only look at patients when they are sick but to be involved in ways of keeping people from getting sick," Lanier said. "By spotting disease markers or indicators of good health ahead of time, we may be able to say why a particular person is susceptible to developing cancer if exposed to a toxin or understanding how we can stop it. Understanding the mechanisms of disease from a predictive modeling perspective and learning how they thrive will provide novel approaches to disease prevention and management."

Utilizing this approach, Dr. Donald DeGracia, an associate professor in the Department of Physiology (WSU School of Medicine), is spearheading the development of a nonlinear mathematical theory of acute cell

injury. The work, being done in collaboration with Dr. Zhi-Feng Huang, an associate professor in the Department of Physics and Astronomy, and Dr. Loren Schwiibert, an associate professor in the Department of Computer Science, will be used to find treatments for such important clinical conditions as stroke, myocardial infarction, traumatic brain injury, and even cancer.

The research stems from the failure of more than 150 clinical trials testing drugs to prevent neuron death after stroke, conducted over the past 30 years and collectively costing billions of

dollars. In contrast to the approach taken by these studies, which placed an emphasis on qualitative biological details, Dr. DeGracia and his team formulated a theory of acute cell injury that uses notions of attractor states and phase planes, well-known in physics but not in biology, to model cell injury.

The theory models a single cell, but since a typical tissue is composed of billions of cells, applying the theory to real cases requires running at least millions of instances of the model simultaneously. Additionally, the team anticipates that interactions amongst the cell instances will greatly increase the required computing resources. While a single instance of the model can be solved in seconds on a modern quad core desktop PC, the parallel implementation over millions of instances requires the state-of-the-art computational resources of the new cluster.

In addition to theoretical considerations, the experimental side of the model requires input data. This will be supplied by Dr. DeGracia's laboratory, which is currently studying gene expression and protein changes after brain ischemia. The results from the lab's microar-

"We don't want to only look at patients when they are sick but to be involved in ways of keeping people from getting sick...Understanding the mechanisms of disease from a predictive modeling perspective and learning how they thrive will provide novel approaches to disease prevention and management."

-Dr. Stephen Lanier, VP of Research

rays and proteomics will be used as data input for solving the model equations. Solving the model requires a computationally expensive pipeline.

The payoff, however, is worth the computational price: in principle, the theory can provide specific predictions for tissue survival or death following certain clinical injuries. In the long term, the work will be used to develop technology that can be used in the hospital to help physicians diagnose and treat various forms of acute cell injury.

“Our work is very appropriate at WSU, which has a large medical school, since the injuries our model addresses are well-known clinical conditions faced by thousands of physicians daily,” said Dr. DeGracia.

“Additionally, WSU has wonderful strengths in physics, computer science, and engineering, which provides a fertile environment for this project to flourish and grow.”

Using mass spectrometry data to find metabolites for cancer diagnostics and drug candidates

Another fertile collaboration is that between Dr. Seongho Kim, an assistant professor at WSU Medical School in the Cancer Biology Graduate Program, and Dr. Hengguang Li, an assistant professor in the Department of Mathematics, who are working on developing parallel computing algorithms for using mass spectrometry (MS) image data to detect biomarkers in terms of metabolites. The field, called metabolomics, is extremely important for cancer diagnostics and cancer drug development.

Using MS data to find metabolites that can be used to develop drugs is a complex problem. Each peak in the data represents one metabolite, but the software that analyzes the data needs improvement to more reliably detect peaks. The more peaks that can be detected, the more metabolites, and thus more possible biomarkers, might be identified.

Working with the MS data, Dr. Kim is trying to develop an improved peak detection algorithm based

on image data, which requires GPU computation. “My work looks at one gigabyte of MS data at a time, which takes a huge amount of time. After joining WSU last year, I looked around for others working with GPUs to collaborate with and found Hengguang Li and contacted him. Our plan is to use the HPC GPU-based cluster to develop a peak detection algorithm. Before, detecting one peak could take about an hour, whereas when the GPU is applied it takes less than 10 minutes.”

Analyzing large genomic datasets to map tissue-specific regulatory sequences

Dr. Roger Pique-Regi, an assistant professor at the Center for Molecular Medicine and Genetics at the WSU School of Medicine, is conducting research that aims to further understanding of the human genome by developing computational methods and statistical models that integrate large datasets, to be used to map tissue-specific regulatory sequences.

A better understanding of the gene regulatory language encoded both in the human genome and in the epigenetic marks that define the state of the cell is a crucial step towards being able to predict the impact of genetic sequence alterations and variation on regulation of gene transcription.

“Learning the grammar that controls the molecular machinery that activates and modulates gene transcription on a given tissue or condition is a fundamental step to be able to understand how the cell is wired and the molecular basis for disease,” said Dr. Pique-Regi.

Developing computational chemistry methods for simulating chemical and biochemical systems

G. Andrés Cisneros, an assistant professor in the Chemistry Department, leads a team that will use the cluster in its work developing and applying computational chemistry methods for the simulation of chemical and biochemical systems. The group contributes to the development of code for the AMBER suite of programs; AMBER is a suite of popular molecular dy-

2014 Research Cluster Grant

namics (MD) software programs for biomolecular simulations and one of the first MD software packages to tap into the power of GPUs.

The group uses computational simulations to investigate several biomolecular systems, including DNA repair and cancer-related proteins. One example, conducted in collaboration with researchers in the chemistry and engineering departments and the Karmanos Cancer Institute, is using the cluster to speed up the development of novel radiolabeled agents for Galectin-3 (Gal-3), a protein involved in such cancers as pancreatic ductal adenocarcinoma. The development of an imaging agent specific for Gal-3 will enable detection of early stage pancreatic ductal adenocarcinoma.

GPU-optimized Monte Carlo simulation engine for molecular systems

Aside from clinical translational research and other similar applications, the cluster's state-of-the-art hardware, particularly the Intel® Xeon Phi™ coprocessors and the NVIDIA® Tesla® GPUs, are being used for projects such as enhancing the development of WSU's novel GPU-optimized Monte Carlo simulation engine, known as GOMC.

The software supports performing high-throughput computational screening of porous materials for CO₂ sequestration, developing novel materials for the stabilization of drug dispersions, predicting phase behavior of polymers and polymer composites, and providing molecular-level insight into such fundamental biological processes as membrane fusion.

Efficient implementation of the Monte Carlo method permits the screening of many potential molecules for desired functionality without any animal or human testing prior to identifying the most promising candidates.

The team, led by Dr. Jeffrey Potoff, a professor in the Department of Chemical Engineering and Materials Science, and Dr. Schwiebert, is developing a general-purpose Gibbs Ensemble Monte Carlo (GEMC)

simulation engine that uses GPUs for acceleration. The GEMC work focuses on atomistic simulation – a computational approach where they determine certain properties of the system by calculating atoms' interactions with each other.

“We've spent a few years developing this new GPU-based Monte Carlo simulation engine to simulate systems that contain up to 2 orders of magnitude more atoms than can be done with other software programs, which is very exciting work indeed,” Dr. Potoff said.

He explained that with the new cluster, the team will be able to model larger systems, reduce data errors, and improve system prototyping. He noted that calculations that might have taken 30 days to complete can now be done in a day.

“Six months of calculations on a CPU take only a week when we add GPUs. It's just that much faster. And this faster calculation time really enhances the research process in a lot of interesting ways.”

New paradigm-shifting methods to solve mathematical models

In addition to collaborating with Dr. Kim regarding metabolomics, Dr. Hengguang Li is pursuing paradigm-shifting ideas in numerical methods for solving mathematical models. As a computational mathematician relatively new to WSU, Dr. Li sought out interdisciplinary research in other departments. He searched for other researchers interested in collaboration and assembled a team that included Dr. Wen Li, an associate professor in the Chemistry Department, and Dr. Seongho Kim, and began talking about collaborating on research.

Numerical simulations are often used in industries where experiments would be expensive or physically impossible. To be accurate, researchers have to be able to look at a large number of smaller pieces that stand in for the domain being modeled. This is “domain decomposition,” the starting point for most numerical simulations; however, breaking the computational domain

into smaller pieces increases the computational complexity exponentially.

To achieve numerical accuracy, a decomposition for a 3-dimensional domain like a cube should consist of billions of small pieces. The computational time for this task alone takes more than a year using an Intel® Core™ i7 3.2 GHz processor, and it is only the first step in the numerical simulation.

Having reached the physical limitation of sequential processing, the next logical step is finding alternative solutions in parallel computing. Dr. Li and his research group are using the new cluster from Silicon Mechanics to develop new parallel algorithms for each module of the finite element method (FEM). This includes algorithms for the domain decomposition, the matrix/vector assembling, and the numerical solver for large systems of equations.

“We are really excited to have the hardware to be able to perform this work,” said Dr. Li. “We are going to be looking at whether different architectures offer different possibilities for working in parallel. We are also looking into whether we can use our existing software with the GPUs, or how we can modify the code to be used on the Intel Xeon Phi coprocessors.”

They hope to incorporate these new algorithms into a software package and plan to make it readily available for computations in various disciplines, in collaboration with colleagues in the chemistry and oncology departments, where numerical algorithms will be used to detect abnormal genes in sequence.

GPU simulation of grain growth and defect dynamics

Another project, led by Dr. Zhi-Feng Huang, an associate professor in the Department of Physics and Astronomy, along with Dr. Schwiebert, currently uses

parallel computation with multiple GPUs to solve and simulate two-dimensional polycrystalline and/or nano-structured material systems.

The work is important in developing new tools used to understand grain growth dynamics, crucial to the control of system properties, such as hardness and strength, of a wide range of engineered materials. The team is extending this to three-dimensional models, for which the use of the new cluster will play an important role.

A great deal of research has been devoted to the understanding of the properties of polycrystalline materials and their design and control. Polycrystalline materials are widespread, from metals to ceramics to minerals. They contain grains of different crystal orientations and grain sizes, shapes, and arrangements, which affect different properties of the materials, including conductivity and fracture.

Production of materials depends on understanding and controlling these polycrystalline structures. The properties of materials arise from their atomic-level, microscopic behavior, but direct study of individual atoms is impractical particularly for large scale systems, so simulations of large-scale material structures are used to gain insights into the materials.

One of the models used to simulate the crystalline materials' behavior is the phase-field crystal (PFC) model, which has enabled researchers to simulate 2D and 3D crystal structures and to study defects. Professors Huang and Schwiebert collaborated on large-scale computer simulations using GPUs to examine various dynamic properties of polycrystalline materials in the PFC model.

“Our goal is to replace programs that can take two months to run on a CPU with much faster turn-around time on the GPU. Currently, we are achieving speedups of about 20-30 times on a GPU for the most compu-

“We could replace a program that runs for 60 days on a CPU with a program that runs in 3 days on the GPU.”

-Dr. Loren Schwiebert,
Associate Professor,
Department of Computer Science

2014 Research Cluster Grant

tationally intensive portions of the code, which means that we could replace a program that runs for 60 days on a CPU with a program that runs in 3 days on the GPU,” said Dr. Schwiebert.

Looking forward towards groundbreaking research

Overall, researchers from WSU have already met with success in using the new Silicon Mechanics HPC cluster to work on complex research, and moving forward, the cluster enables researchers to dream bigger and plan projects that would have seemed impossibly complex or time-consuming a short time ago.

“Analyses are often limited by the computing resources that are available,” said Gossman. “Many analyses that may have taken years to wade through in the past, can now be processed in weeks or even days on the Silicon Mechanics cluster, greatly accelerating the speed of discovery. The new cluster donated by Silicon

Mechanics is an important new tool in our advanced computing environment and one that gives our researchers the necessary processing speed to tackle even larger, computation-intensive problems in healthcare, bioinformatics and genomics, as well as in advanced manufacturing and materials science.”

“Many analyses that may have taken years to wade through in the past, can now be processed in weeks or even days on the Silicon Mechanics cluster, greatly accelerating the speed of discovery.”

-Patrick Gossman, Deputy CIO,
Community, Research
and Special Projects

About Silicon Mechanics

Silicon Mechanics, Inc. is an industry-leading provider of rackmount server, storage, and high-performance computing solutions. Deploying the latest innovations in hardware and software technology, we work in collaboration with our customers to design and build the most efficient, cost-effective technology solution for their needs. Our guiding principle, “Expert included,” is our promise that reflects our passion for complete customer satisfaction, from server and component selection to superior installation and ongoing technical support.

Silicon Mechanics has been recognized as one of the fastest growing companies in the Greater Seattle Technology Corridor.

About Wayne State University

Wayne State University is a metropolitan research institution dedicated to preparing students to excel. It does so by combining the academic excellence of a major research university with the practical experience of an institution that by its history, location and diversity represents a microcosm of the world we live in. As a nationally recognized urban public research university, Wayne State’s mission is to create knowledge and prepare a diverse body of students to excel in an increasingly complex and global society.

Founded in 1868, the university offers more than 370 academic programs through 13 schools and colleges to nearly 28,000 students. Wayne State’s main campus in Midtown Detroit

comprises 100 buildings over 200 acres; its six extension centers offer higher education to students throughout Southeast Michigan.

About Art Mann

Art Mann is a software sales and business development professional with more than 25 years of experience serving the educational and research market. He has served as Silicon Mechanics’ education/research /government vertical group manager for the past 10 years, focusing on strategic planning and developing and growing new and installed base opportunities. Mr. Mann’s long career supporting educational and research needs gives him a special expertise.

Prior to Silicon Mechanics, Mr. Mann founded Mann Technology Advisors, an information systems solutions company specializing in the higher education market, which was successfully sold to Denali Advanced Integration in 2003. He previously worked for Compaq Computer Corporation/Digital Equipment Corporation (now HP) for 17 years. Mr. Mann has served as a lecturer in the Math-Sciences Division and assistant dean of the Graduate School at Babson College, and as a lecturer at the University of Washington’s Graduate School of Business. He has an MBA from the University of Massachusetts at Amherst, and a BS in Management from State University of New York at Buffalo.



Dr. Stephen Lanier



Patrick Gossman



Dr. Loren Schweibert



Dr. Hengguang Li



Dr. Jeffrey Potoff



G. Andrés Cisneros



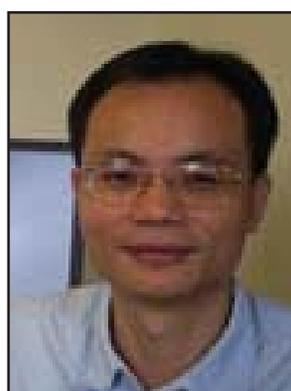
Dr. Wen Li



Dr. Roger Pique-Regi



Dr. Seongho Kim



Dr. Zhi-Feng Huang

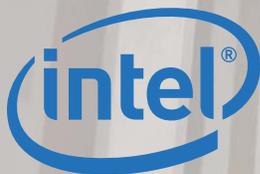


Dr. Donald DeGracia

Case Study made possible by:



and our partners:



NVIDIA

