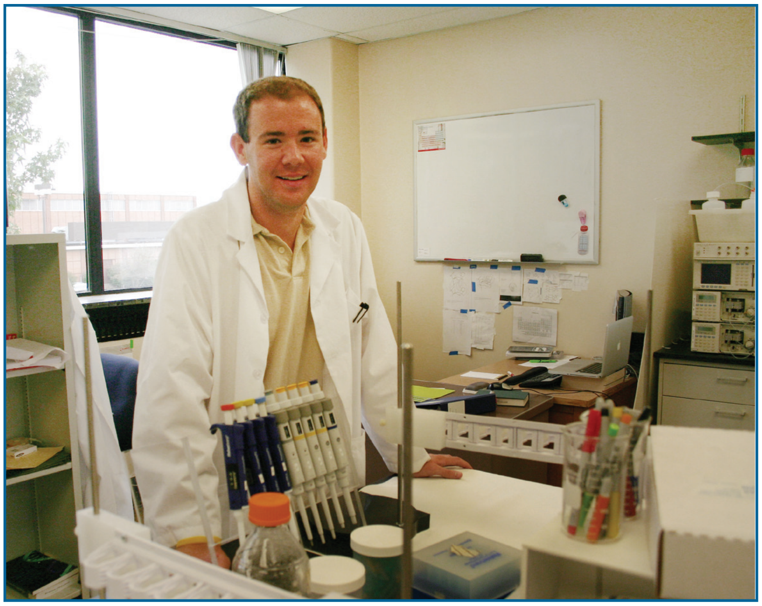

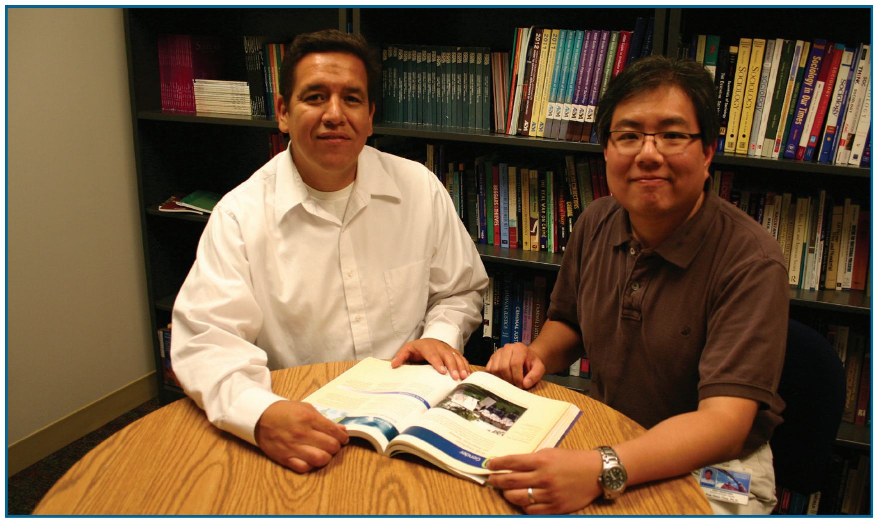

SAINT LOUIS UNIVERSITY

Research Cluster Grant Update

High-Performance Computing Cluster Fostering Great Research Results at SLU

High-performance computing cluster fostering great research results at Saint Louis University

By Art Mann, Silicon Mechanics

We recently checked in with Saint Louis University (SLU), of St. Louis, Missouri, to see how their research efforts were advanced by the complete high-performance computing (HPC) cluster awarded by Silicon Mechanics as part of our Research Cluster Grant program.

What we found was an incredibly animated group of researchers, thrilled to be able to advance their varied research interests, and convinced that the increased computational firepower is absolutely central to doing their cutting-edge research better and faster.

The compute cluster was developed by Silicon Mechanics, and includes hardware donated by Kingston Technologies, AMD, NVIDIA, QLogic, Supermicro, and Seagate, with cluster management software supplied by Bright Computing.

University's technology leader uses cluster to advance research goals

Keith Hacke, SLU's chief technology officer, is on a quest to help the university reach its goal of becoming a Top-50 research institution. ***"Having high-performance computer clusters has been very important to us on a number of fronts.*** Not only does it help faculty do research, but it also serves as a recruiting and publicity tool, helping us to accomplish our research mission and fill our faculty ranks."

Hacke explains that the addition of the new HPC cluster has been an eye opener, and he especially cites the cluster management solution from Bright Computing as a boon to SLU's overall computing efforts. "We have adopted the software as our standard cluster operating system, which has provided us a great deal of capabilities. It makes it far easier to run the cluster and expand it beyond the original equipment."

Keith Hacke

He says the GPU (graphics processing unit) accelerators have also turned out to be a wonderful resource, with more benefits to come. "We are still learning how to use the GPU to its best advantage. Applying GPUs to actual real-world problems requires some pretty advanced skills. But we have been growing those skills here at the campus, so it has definitely expanded our horizons."

Departments sharing the research cluster include Aerospace and Mechanical Engineering, Biology, Center for Digital Theology, Chemistry, Sociology (jointly with Earth and Atmospheric Sciences, Center for Sustainability, and School of Public Health), and Political Science, as well as the John Cook School of Business.

HPC cluster report card – making great progress!

SLU was awarded the grant from a pool of applicants comprised of US and Canadian universities and research institutions. The application from Saint Louis outlined a unique multidisciplinary resource-sharing approach: the cluster is being used for seven different research initiatives, encompassing ten university departments.

Several leaders of the research projects using the cluster gave us an update on their progress, focusing especially to how they are using the HPC cluster – and sharing a few unanticipated or unexpected benefits.

Table 1 summarizes the teams working with the cluster, their original research goals, and a brief overview of their progress thus far.

First to report is the **Aerospace and Mechanical Engineering Department**, which is using the cluster for a variety of computational research projects, including airfoil studies, parachute studies, airdrop system research, and parachute plus payload system research for the Army, as well as some computational research for Boeing. The research is led by Dr. Mark McQuilling, with collaboration from parachute inflation expert Dr. Jean Potvin, of SLU's Physics Department.

Mark McQuilling

Airdropping Humvees

In research funded largely by the US Army Natick Soldier Research, Development, and Engineering Center (NSRDEC), one of the key concepts being studied is an annular parachute, because annular geometries hold the potential for the highest achievable drag.

Dr. McQuilling reports that his research on new parachute concepts may one day enhance the ability to accurately deliver cargo by airdrop. And by cargo, think big – like delivering a Humvee on a platform to soldiers in remote areas!

"If you can imagine a parachute that looks like a hemisphere of a cup, if you cut holes in it, in just the right places, it produces more drag, which is what you're looking for," said McQuilling. "But the problem with cutting holes in the parachute is it can get pretty unstable, and the unstable nature of that parachute has prevented its widespread use so far." Dr. McQuilling and Dr. Potvin have been studying the annular parachute to make it more stable. "If we can figure out this stability issue, we can develop a low-cost alternative to the existing methods being used for dropping cargo."

Calculations stepped up

Previously the researchers had been using a cluster funded by the US Army NSRDEC group, but the new cluster has allowed them to expand their simulation capability to more airdrop system geometries. "We are studying the typical hemisphere air parachutes, cargo personnel, and payload, and we are just about ready to simulate the Humvee platform."

Dr. McQuilling explains that without the cluster, the types of simulations his team works on would not be possible. Using the Humvee geometry as an example, ***one simulation with one angle of attack might take three to six months on a single machine, but only a matter of weeks with a cluster.*** "So we can get more angles of attack, and more resolution of the different physics occurring around the real thing that is being dropped."

Studying problems caused by drop in air density

In addition to the Army work, Dr. McQuilling is studying low-pressure jet engine turbine aerodynamics, with the goal of improving turbine airfoil performance. "The research has allowed us to figure out how to let these jets fly higher when air density drops, which means there is not as much air for the combustion process. The HPC cluster lets us study the problems this causes in the air turbine sections. This research will ultimately be used to develop engines that are smaller, lighter, weigh less, perform better, and are able to fly higher."

The **Biology Department** is doing research on climate change, using the research cluster to model high-elevation plant community ecology in the Andes Mountains. Dr. Gerardo Camilo explains that SLU researchers conducted exhaustive research on a group of plants that grow at high elevations in the Andes, gathered by viewing and documenting collections in South American universities, as well as trekking through the mountains themselves.

Researchers are gathering as much information as possible on the present and past distribution of all these species, so they can generate a confirmed trajectory, allowing them to make projections into the future.

The urgency of climate study

Dr. Camilo explains that his team will also be using the cluster to compare different scenarios from the Intergovernmental Panel on Climate Change (IPCC), looking at projections for 2020, 2040, and 2060. The IPCC just released their latest assessment report, indicating that worldwide temperatures are probably going to rise almost 4 degrees Celsius rather than the two and a half degrees previously predicted. "This puts greater urgency into the need to tackle climate change," said Dr. Camilo, who expects to start his work in the spring of 2014.

Another interesting use for the HPC cluster is on related research, exploring what happens to climate change in the face of multi-use conservation effort at very large scales. Rather than the U.S. approach to conservation, which focuses on minimizing human impacts, European efforts tend to focus more on the interactions between how people utilize land and how they carry out their cultivation efforts. Dr. Camilo hopes to apply the research cluster to running different hypothetical conservation scenarios, studying which combinations of factors would help preserve biodiversity.

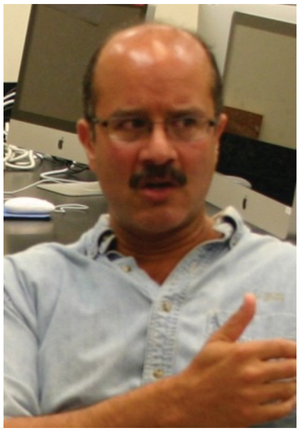

Gerardo Camilo

The age of big data

Dr. Camilo also considers the HPC cluster to be invaluable for training his students to be scientists in the 21st century. "We have entered the age of big data. ***The idea that somehow you can do your research in your own little lab, that's so 20th century. We need to train our students with a broader set of skills***."

He explains that students need to learn how to retrieve a data set, manipulate the data, get it to the appropriate stage for the analysis they are going to perform, and analyze the data and try to interpret results. As part of this effort, he teaches a course in computational applications in biology, training graduate students to develop their own projects, search for data, and perform their own analyses.

Modeling requires significant compute power

He adds, "Laptop computers are great, and 99 percent of what we do is done on a laptop, including scripting and debugging. But when you get into serious modeling, you must move to the HPC environment. ***With HPC you can actually start asking questions that ten years ago nobody would have ever thought could be answered.*** At the same time there is also a real sense of urgency, that we need to get this done really fast."

Chemistry Department Associate Professor Chuck Kirkpatrick, PhD, discussed his research into carbon nanotubes, which are being investigated for everything from nano-wires, to artificial muscle tissues, to textiles. Carbon nanotubes tend to cluster together, and Dr. Lewis is using the cluster to develop a list of possible solvents that can be used to separate them.

Chuck Kirkpatrick

"The research that we do is focused on weak interactions between molecules," said Dr. Kirkpatrick. "Chemists are pretty good at computing strong interactions, but weak interactions are very difficult and that's where the cutting edge of computational chemistry is these days. For this project, we are looking at how solid molecules would interact with carbon nanotubes."

Separating natural attraction

Scientists are looking at how carbon nanotubes can act as catalysts and also how they can be attached to other molecules so they can be used for electro-chemical measurements. Experimentalists working in the field have a real challenge because the carbon nanotubes need to be separated to be useful.

"When we put them into a solvent, we want the solvent to separate them and most solvents don't do that. So our goal is to be able to calculate the interaction or the strength of the interaction between the solvent and the nanotube. Once we come up with a shortlist of candidate solvents, experiments could begin."

The team takes small parts of the molecule and does individual calculations on its various regions. The goal is to narrow the field from thousands of potential solvents to around 20 candidates. With the many CPU cores in the new high-performance computing cluster, Dr. Kirkpatrick can run calculations in parallel to dramatically speed up the computational process.

Using GPUs to make molecule movies

The chemistry researchers are looking forward to better exploiting the GPUs in the HPC cluster. "What I would like to be able to do is view the nanotube as a solvent molecule approaches it, look at how the nanotube and solvent molecule interact, and literally make a movie about the interactions," said Dr. Kirkpatrick. ***He intends to incorporate visualization of molecular interaction into the department's computational chemistry courses.***

Dr. Justin Goodson, from the **John Cook School of Business** Department of

Operations and Information Technology Management (ITM), is using the cluster to help solve large-scale optimization problems. He explains that the research will be used in transportation and logistics, specifically the routing of delivery vehicles, including anything from local milk delivery trucks, to less than truckload (LTL) shipping companies.

Transportation uncertainty

According to Dr. Goodson, for the past 60 years dynamic programming has provided a useful framework for modeling complex problems. But because of the uncertainty involved, finding optimal solutions to a dynamic program can be quite challenging.

In some of the transportation and logistics problems he focuses on, the uncertain factor is demand. For example, when a shipping provider contracts to pick up four pallets, if the customer has six pallets when the truck arrives, the vehicle may run out of capacity. This type of last-minute change and uncertainty in customer demand makes it very difficult to plan efficient delivery routes.

Optimal approximations

Justin Goodson

Dr. Goodson asserts that it is computationally impossible to completely solve some of these problems. Finding an exact solution would require more computer memory than there is available on the earth, and more time than from the Big Bang till now! "My research focuses on coming up with approximate solutions that seem to work well in practice and are perhaps within a certain distance of what we believe provides an optimal solution that will be of value to people."

Dr. Goodson's work uses the research cluster to model these problems as stochastic dynamic programs. Since it is difficult to find an exact solution, he has worked on finding approximate solutions by leveraging the large number of processors to perform thousands of experiments to determine how these algorithms perform, based on a set of test incidences.

Efficient transportation saves money

Dr. Goodson points out the practical benefits of this research. ***"Transportation is becoming an increasingly large component of supply chain cost, so any reductions to transportation costs go immediately to your bottom line."***

The main benefit of the computing cluster is the significant speedup of research time. Now that he has better access to a cluster, Dr. Goodson reports that he can run his experiments much faster. "Even with the large amount of processors that we have, running a full set of experiments can sometimes take a few weeks. Using the high-performance computers, I can perform large numbers of experiments in quick succession."

New reality makes new solutions possible

He also adds that high-performance computing is going to play an increasing role in how we design algorithms to solve computationally demanding problems, so the HPC cluster is giving us a glimpse at what we can do. "***Ten years ago many people would have said you couldn't possibly solve a problem like that and now it's a reality.*** I wonder what I will be doing with these computers ten years from now."

One fascinating project that will make use of the HPC cluster is jointly sponsored by the **Sociology & Anthropology Department**, in cooperation with Earth and Atmospheric Sciences, Center for Sustainability, and the School of Public Health (Department of Environmental and Occupational Health).

J.S. Onésimo (Ness) Sandoval, PhD, an associate professor in the Department of Sociology and Anthropology, talked about his research on a new kind of methodology for synthesizing information on the spatial aspects of social and economic, environmental and ecological phenomena. The project is known as Mapping Risk and Resilience in Saint Louis (MaRRS).

Ness Sandoval

Many measurements, one indicator

The project was motivated by a call from the National Science Foundation for researchers to come up with a methodology for social and environmental synthesis. Rather than bringing together people from different disciplines who worked only on their area of expertise, they would work together as a holistic team. The goal would be to validate a methodology to synthesize all of the information they collected to come up with a new understanding of the social phenomenon they were interested in.

"In this case we were interested in inequality or risk, so we were looking at social and environmental synthesis, bringing together different types of data that have different levels of measurements and different units of analysis, to come up with one single indicator."

St. Louis testing ground

The city of St. Louis itself is the first testing ground for the method, which will be replicated in five other cities. As a city that has undergone tremendous change from 1950 through 2010, St. Louis was a good laboratory in which to frame the new social-environmental synthesis methodology around the theoretical question, "What happens to a city when it loses two thirds of its population? Is risk randomly distributed throughout the space, or does it become concentrated in certain parts of the city?"

The team looked at four dimensions of risk: social, economic, environmental, and ecological. Ultimately, they collected 18 dimensions of social data, 9 dimensions of economic data, 7 dimensions of environmental data, and 5 dimensions of ecological data. All the ratings were on different scales, using different units.

Rather than measure these dimensions using census blocks or other large governmental units, they created a grid of the city made up of 250-meter by 250-meter squares. They developed a formula to standardize the information to translate it into the grids, and then rated the risk for each. The social environmental synthesis risk index was created from these different dimensions of inequality.

Originally, the team had tried to analyze all the data on a single computer, which took "forever." Reaching out to SLU's information technology department, they learned of the research cluster, and jumped at the chance to have access to additional computational resources for the massive set of spatial data.

Training students in a changing environment

As part of their grant, the team trained numerous students on "R" software for data manipulation, calculation, and graphical display. They are now improving the efficiency of their analyses using SAS statistical analysis software on the cluster. "I would say 95 percent of our innovation is spatial statistics that is taking place in R, so you need massive power and massive space, which makes the HPC cluster the perfect environment to do these big comparisons of cities we hope to do in the next phase of the research.

"We are from the social sciences, where we are used to using our desktop computers, so I think a real culture change is now happening."

We are inspired by research being conducted at SLU

Silicon Mechanics developed the Research Cluster Grant competition because we believe that putting computational resources in the hands of researchers can change the way science is practiced. The research at Saint Louis University has embodied this objective.

We have been inspired by the diversity of cutting-edge research being transformed by access to the HPC cluster. The university has benefited from the visibility the award has generated. Perhaps best of all is the opportunity for students to participate in the real-world learning opportunities enabled through these important projects.

Table 1

High-Performance Cluster Research Teams – 2013 Update

We checked in with several of the departments benefitting from the first Research Cluster Grant to find out how their work is progressing, and learn more about how these researchers are able to use the cluster. Here are a few highlights of the interviews.

Department	Research Goals	Researchers Interviewed	How the HPC Cluster is Enhancing Research
Aerospace and Mechanical Engineering Department, Parks College of Engineering, Aviation and Technology	Study new parachute concepts to enhance the ability to accurately deliver cargo by airdrop.	Mark McQuilling, PhD	Exploring stable annular parachute geometries, for airdropping cargo like Humvees. The cluster allows many more simulations, in less time. For example, a simulation that takes 3 - 6 months on a desktop computer takes only a few weeks on a cluster.
Biology	Expanding into the areas of climate change, modeling high-elevation plant community ecology in the Andes Mountains.	Gerardo Camilo, PhD	Gathering and analyzing information on the distribution of species of interest in the Andes, using past and present data to project future populations. The cluster will be used to analyze the vast amount of data, and run hypothetical scenarios to assess various conservation efforts. The research cluster is also key to training students to use modern hardware to grapple with big data problems in their own research.
Chemistry and Engineering Department	Studying carbon nanotubes and developing computational models to predict solubility in organic solvents.	Charles C. Kirkpatrick, PhD	Calculating the interactions between carbon nanotubes and thousands of possible solvents, developing a short list of solvents for future research. The many processing cores in the research cluster speed up calculations through parallel computations.
John Cook School of Business, Department of Decision Sciences and Information Technology Management (ITM)	Developing algorithms to solve large-scale stochastic dynamic programs with the goal of improving logistical decision making.	Justin Goodson, PhD	Modelling large-scale optimization problems in transportation and logistics. The calculations have so many variables that researchers use the research cluster to run thousands of tests on each algorithm, looking for good approximations of real-life routing and capacity problems.
Sociology & Anthropology **School of Public Health** **Earth and Atmospheric Sciences** **Center for Sustainability**	Mapping Risk and Resilience in St. Louis (MaRRS).	J.S. Onésimo Sandoval, PhD	Developing a new method of synthesizing social data, bringing together different measurements with different units of analysis, to create a single meaningful indicator of social and economic risk. The HPC cluster provides the compute power necessary to manipulate the massive set of spatial data. The department is able to train students to develop and improve their analyses using statistical analysis software on the research cluster.

Bio

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 eight 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, and he recently initiated a series of technology-focused seminars for research faculty and their information technology support.

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 Saint Louis University

Saint Louis University is a Catholic, Jesuit institution that values academic excellence, life-changing research, compassionate health care, and a strong commitment to faith and service. Founded in 1818, the University fosters the intellectual and character development of nearly 14,000 students on two campuses in St. Louis and Madrid, Spain. Building on a legacy of nearly 200 years, Saint Louis University continues to move forward with an unwavering commitment to a higher purpose, a greater good.

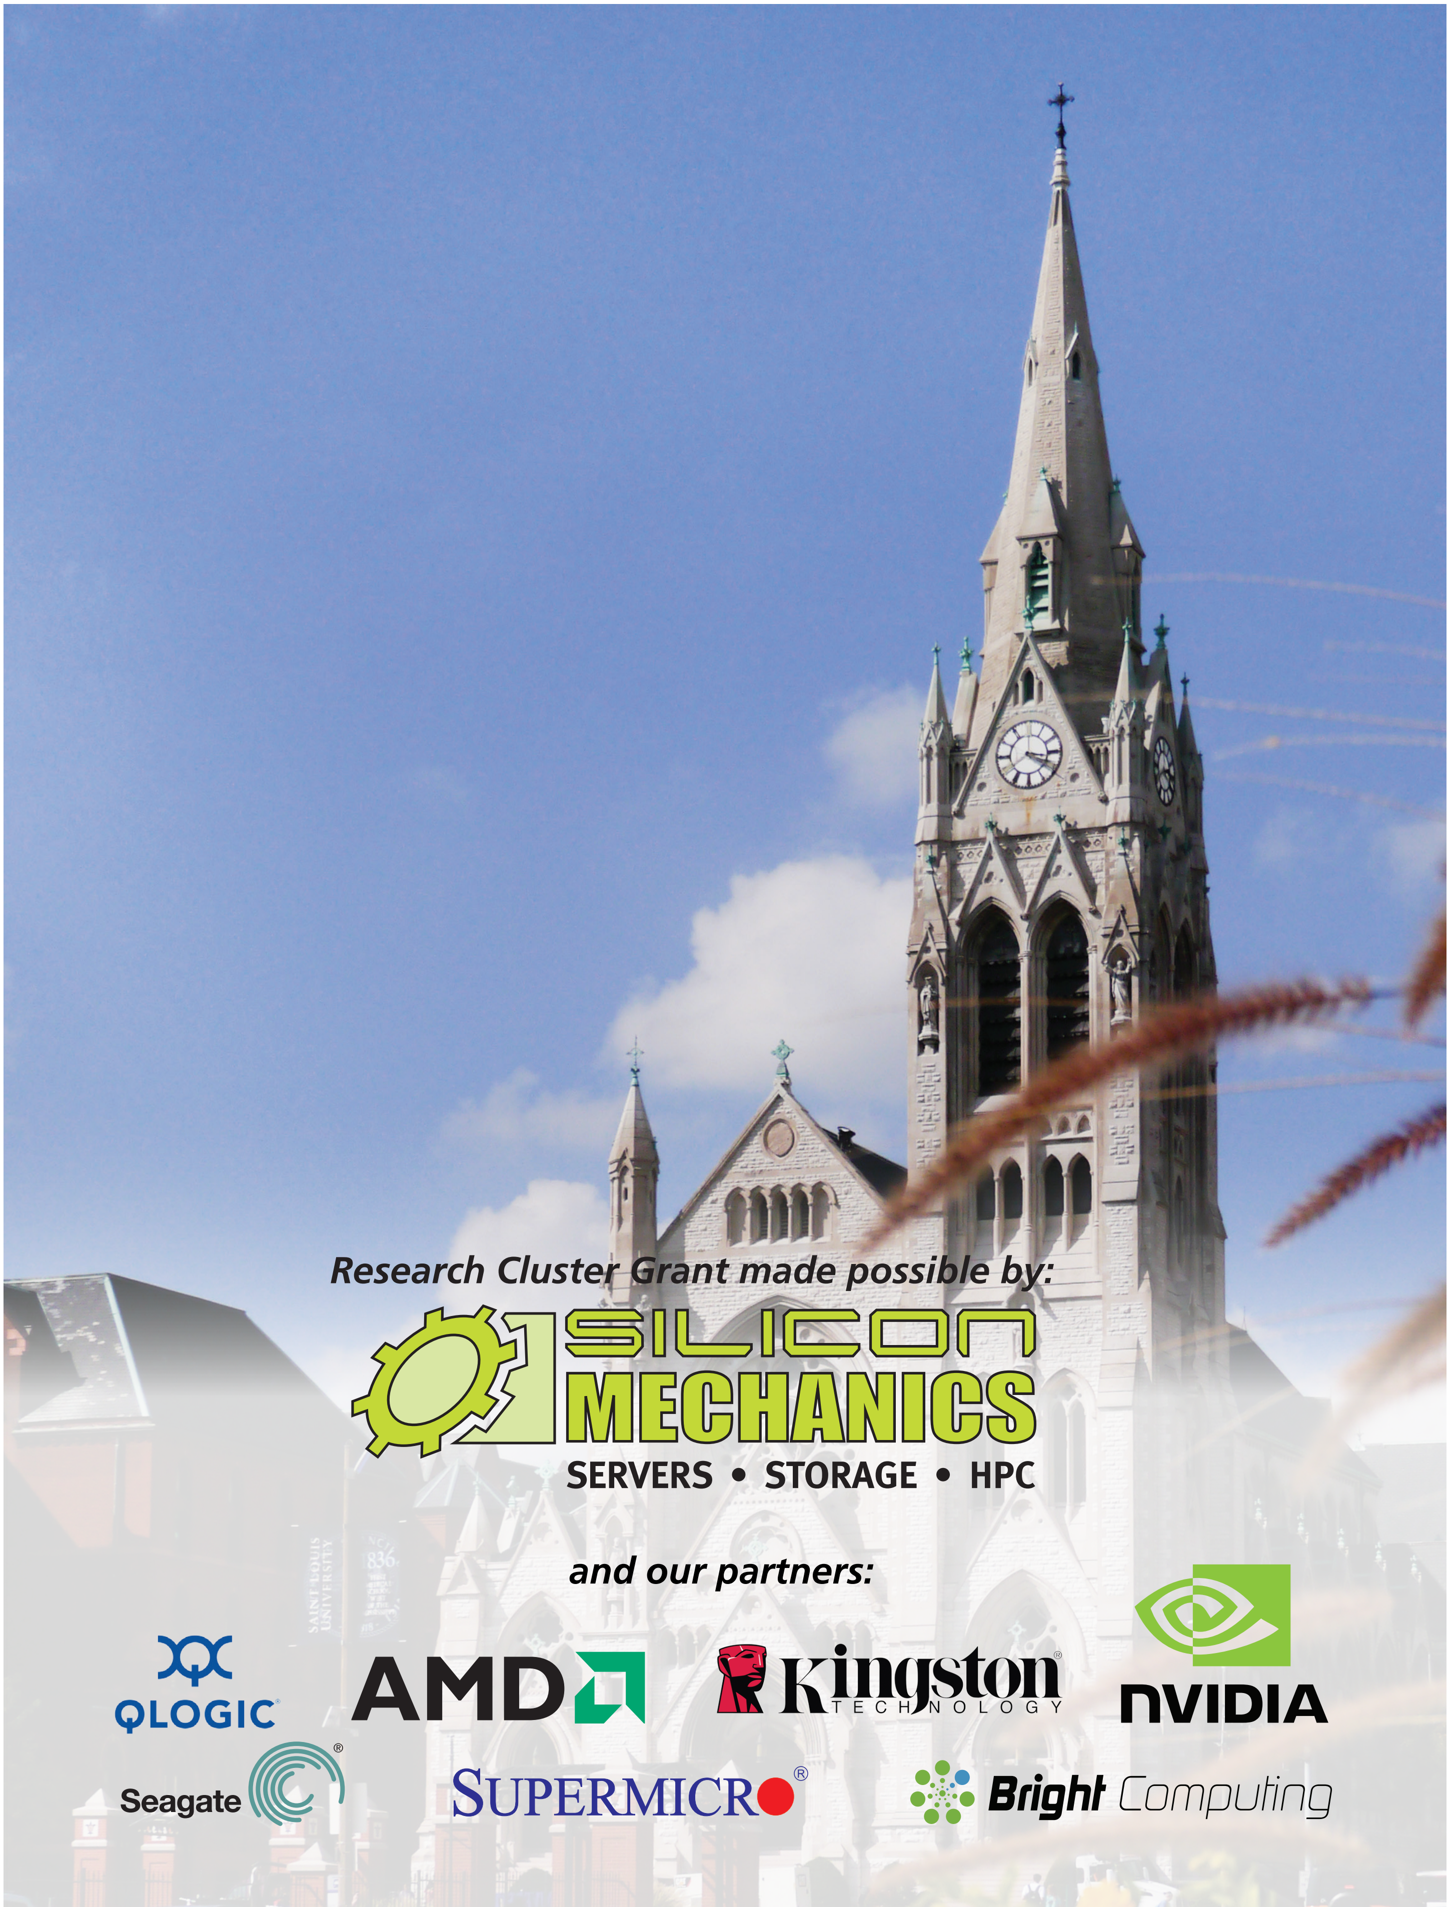

Research Cluster Grant made possible by:

SILICON MECHANICS

SERVERS • STORAGE • HPC

and our partners:

QLOGIC

AMD

Kingston TECHNOLOGY

NVIDIA

Seagate

SUPERMICRO

Bright Computing