



Business challenge

U-M wanted to accelerate scientific discovery in fields such as cosmology and medicine. The problems involved are so complex that even the latest supercomputers struggle to accurately model them.

Transformation

Modeling the real world accurately enough to make meaningful predictions is extremely difficult, even with continuous growth in computational power. By enhancing predictive modeling techniques with the combined power of supercomputing, AI, and machine learning, The University of Michigan (U-M) and IBM are helping scientists to make new discoveries.



Karthik Duraisamy
Associate Professor
at the Department of
Aerospace Engineering
and Director of the
Center for Data-Driven
Computational Physics,
University of Michigan

Business benefits

Enables

doctors to perform surgery with unparalleled insight and precision

Helps

improve the design of aircraft and rocket engines

Deepens

our understanding of the expansion of the universe

University of Michigan

Accelerating the pace of ground-breaking scientific innovation with HPC and AI

One of the foremost research universities in the United States, the [University of Michigan](#) offers a comprehensive range of doctoral and professional degrees. Founded in 1817, it is the state's oldest university, and has more than 27,000 undergraduate and 13,000 graduate students.

“IBM was an ideal choice for both servers and storage. In our benchmarks, the IBM Power System S822LC performed better than the competition.”

—Karthik Duraisamy, Associate Professor at the Department of Aerospace Engineering and Director of the Center for Data-Driven Computational Physics, University of Michigan

Share this



Uniting big data and supercomputing

Scientific research involving physics-based models is enormously complex. For example, predicting the performance of an aircraft engine involves modeling innumerable molecular interactions every second. For nonlinear problems like these, the extreme sensitivity to starting conditions—where a minor change can quickly produce a dramatic difference in outcomes—means that even the latest supercomputers can only produce approximate results.

Traditionally, fluid dynamics and other nonlinear modeling has been carried out using high-performance computing (HPC) environments, while big data analytics has largely been confined to data-parallel architectures. Scientists from the University of Michigan (U-M) realized that they could massively improve the accuracy and reliability of their predictive models by bringing these two worlds together. Their aim was to combine enormous volumes of observational and experimental data with real-time machine learning—delivering more accurate simulation results.

With these two technologies combined, researchers would be able to drill down into the details of their simulations by integrating data from other experiments into their models—eliminating the need to run additional time-consuming and costly simulations. Researchers recognized that this could accelerate the pace of innovation and change the face of science across multiple disciplines.

Karthik Duraisamy, Associate Professor at the Department of Aerospace Engineering, and Director of the Center for Data-Driven Computational Physics at the University of Michigan, explains: “There was a pressing need for data-driven predictive modeling to help re-envision traditional computing models. So, we decided to build ConFlux, an entirely new resource to enable traditional HPC applications to interact with big data and machine learning.”

Model for success

The university’s first step towards making ConFlux a reality was choosing a vendor for its HPC infrastructure.

Duraisamy says: “IBM was an ideal choice for both servers and storage. In our benchmarks, the IBM Power System S822LC performed better than the competition. Its low-latency architecture, IBM POWER8 processors, and integrated NVIDIA Tesla GPUs offered better performance than equivalent x86-based systems. To store and manage our big data, we needed a high-performance storage solution that we could seamlessly scale out as data volumes grew, and IBM Elastic Storage Server (ESS)—based on IBM Spectrum Scale—was the perfect fit.”

In addition to offering eight hardware threads per core and more on-chip cache than x86 processors, IBM® POWER8® features IBM’s Coherent Accelerator Processor Interface (CAPI) technology. CAPI provides a high-performance solution for customizable, computation-intensive algorithms. Where a traditional accelerator can access shared memory only via the main processor’s I/O subsystem, CAPI connects acceleration engines directly to the coherent fabric of the POWER8 chip, removing complexity, cutting latency, and resulting in significant overall performance gains for many applications.

For GPU accelerated workloads, IBM Power System S822LC for HPC servers are unlocking new performance from massively parallel and tightly connected Tesla P100 with NVLink GPUs. The advanced Phase 2 POWER8 with NVLink nodes in ConFlux feature a broader pipe delivering 2.5 times the bandwidth CPU:GPU of alternatives, for superior application performance.

By incorporating Mellanox InfiniBand fabric, the IBM HPC platform delivers high bandwidth and low latency when jobs span the cluster. This means that the IBM POWER8 processors spend more time doing useful work and less time waiting for data. With physics simulations that can last days, or even weeks, the cumulative impact of eliminating processor idling represents a huge performance advantage.

IBM ESS combines the parallel file system of the IBM Spectrum Scale™ (formerly GPFS) solution with IBM POWER8 servers and dual-ported storage enclosures.

IBM Spectrum Scale facilitates system throughput growth while still provisioning a single namespace, eliminating data silos, simplifying storage management, and delivering peak performance. By consolidating storage requirements across an organization, IBM ESS helps reduce inefficiency, lower acquisition costs, and support demanding workloads.

Duraisamy adds: “With the Power Systems approach, IBM is bringing together leading technologies from the OpenPOWER foundation — NVIDIA GPUs, Mellanox InfiniBand networking, and POWER8 processors—in a single HPC platform.”

Todd Raeker, Research Technology Consultant for the University of Michigan, comments: “The workload on ConFlux will be hugely data-intensive. Any given simulation will generate 10 to 20 terabytes of data, and there could be dozens of simulations running at one time. With this in mind, we expect to see surges in storage requirements as the datasets grow — which is why we chose

IBM Spectrum Scale as our software-defined storage solution. With Spectrum Scale, we can seamlessly increase performance and capacity by simply adding nodes.”

Raeker expands on the advantages of choosing one vendor for both computing and storage: “One of our objectives was to build a relationship with a single vendor, so rather than having separate vendors for software and hardware, we wanted IBM to be responsible for the whole stack. We knew that with everyone on the same page, we would be able to resolve potential issues much more quickly, avoiding downtime.”

The university will be running a variety of big-data-driven workflows, including some on Apache Hadoop. With IBM Spectrum Scale, U-M can effectively layer a Hadoop Distributed File System (HDFS) type environment over a Spectrum Scale environment, greatly increasing flexibility.

Raeker says: “We can work with data in HDFS and Spectrum Scale simultaneously, which significantly simplifies scientists’ workflow; they

can focus on their experiments instead of worrying about whether the infrastructure is in the correct state to support them.”

With an environment optimized for rapidly moving large volumes of data into and out of compute nodes, U-M can ensure that results from data queries are delivered much faster than before. Crucially, machine learning algorithms will then pull fresh insights from these results and use them to adapt predictive models in real time, continuously improving the accuracy of simulations even as they are in progress.

For the first phase of the deployment for ConFlux, U-M installed IBM Power System S822LC with POWER8 processors alongside an IBM Elastic Storage™ Server, and in the second phase it deployed additional compute nodes with newer POWER8 with NVLink processors and NVIDIA Tesla P100 with NVLink GPUs.

Duraisamy comments: “IBM explained to us that we could get an S822LC with POWER8 straightaway, and then move up to

the next generation of processors when they become available. It was great knowing that we could invest immediately and also take advantage of future performance enhancements.”

Unlocking innovation

With ConFlux in place, the university is set to advance the threshold of scientific knowledge in a wide range of fields. Already, the HPC at U-M is being used for several projects, including studies in cardiovascular disease, turbulence, and dark matter.

For example, a strong sign of certain cardiovascular diseases like hypertension is arterial stiffness. By combining noninvasive imaging techniques, such as MRI and CT scans, with a model of blood flow produced by ConFlux, scientists at the university hope to be able to accurately estimate arterial stiffness in less than an hour, helping doctors to deliver the appropriate treatment to patients more quickly than ever.

Even though this research is in its early stages, scientists in this field are already realizing the benefits of the ConFlux HPC environment, with one academic noting in a [recent journal](#) that this innovative approach could save significant amounts of time in the initial assessment of cardiovascular diseases, and improve outcomes by helping establish the best way to approach surgery in each case.

When it comes to designing efficient aircraft and rocket engines, understanding turbulence is key, yet the patterns that air produces when it breaks up are exceptionally difficult to predict. U-M is working with leading engineers to improve the [accuracy of turbulence simulations](#) using ConFlux, enabling better testing and accelerating the development of superior aircraft designs.

The university also intends to use ConFlux to better understand our universe. By feeding data from large galaxy-mapping studies like [The Dark Energy Survey](#) into simulations of galaxy formation, U-M hopes to gain a better understanding of the role that dark matter plays in the continuing expansion of the universe—potentially driving a revolution in the understanding of core scientific concepts such as gravity.

Raeker adds: “With IBM hardware boosting our HPC environment, we can offer scientists the tools to conduct research that could revolutionize entire industries. Because Conflux is at the bleeding edge of supercomputing technology, we are continually learning how to optimize our architecture. IBM hardware provides us with the flexibility to adapt to a variety of experimental scenarios so we can always create the right conditions for innovation.”

Looking further to the future, U-M has plans to share the new platform with research groups at other universities, and is looking to integrate material learned from ConFlux into its degree programs.

Duraisamy concludes: “ConFlux allows us to bring together large-scale scientific computing and machine learning to accomplish research that was previously impractical, given time and resource constraints. This capability will close a gap in the US research computing infrastructure and place U-M at the forefront of the emerging field of big data-driven physics.”

Solution components

- IBM® Elastic Storage™
- IBM PowerLinux™
- IBM Power System S822LC
- IBM Power System S822LC for HPC
- IBM Spectrum Scale™
- Mellanox InfiniBand

Connect with us



Take the next step

To learn more about IBM Power and IBM Storage solutions, please contact your IBM representative or IBM Business Partner, or visit the following websites: ibm.com/power and ibm.com/storage.

© Copyright IBM Corporation 2017. 1 New Orchard Road, Armonk, New York 10504-1722 United States.

Produced in the United States of America, October 2017. IBM, the IBM logo, ibm.com, Elastic Storage, IBM Spectrum Scale, PowerLinux, and POWER8 are trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the web at “Copyright and trademark information” at ibm.com/legal/copytrade.shtml. This document is current as of the initial date of publication and may be changed by IBM at any time. Not all offerings are available in every country in which IBM operates. THE INFORMATION IN THIS DOCUMENT IS PROVIDED “AS IS” WITHOUT ANY WARRANTY, EXPRESS OR IMPLIED, INCLUDING WITHOUT ANY WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND ANY WARRANTY OR CONDITION OF NON-INFRINGEMENT. IBM products are warranted according to the terms and conditions of the agreements under which they are provided.



POC03334-USEN-00

