HYPERION RESEARCH

HPC User Forum Update

HPC in Industry, Brendan McGinty, NCSA Industry Director

Bob Sorensen, Melissa Riddle
February 2020

IN THIS UPDATE

The HPC User Forum was established in 1999 to promote the health of the global HPC industry and address issues of common concern to users. In September 2019, the 73rd HPC User Forum took place at Argonne National Laboratory in Illinois. This update summarizes a presentation from that meeting entitled *HPC in Industry*, given by Brendan McGinty, NCSA Director.

"For over 30 years, NCSA Industry has helped innovative organizations overcome their biggest obstacles. Today, we partner with government groups and global companies to tackle big questions and provide meaningful results in fields ranging from artificial intelligence to genome mapping to autonomous transportation–all supported by our eight-time HPCWire Award-winning team of experts and delivered on advanced computing resources."

Source: NCSA and Hyperion Research, 2020

HPC IN INDUSTRY, BRENDAN MCGINTY, NCSA DIRECTOR

Brendan McGinty presented an overview of the NCSA Industry Program including its history, goals, and case studies with demonstrative ROI. He began his presentation with a short video showcasing some of the early partners in the NCSA Industry Program including Caterpillar, Boeing, Eli Lilly, FMC Corporation, Motorola, JP Morgan, and Allstate. The video showcased how the NCSA Industry team works across domains to deliver on the promise of HPC.

McGinty stressed that maintaining an industrial culture and pace is an important aspect to the team's ongoing success. "*Unlike academia, industry success requires quicker hits and quicker deadlines. These milestones mean everything to companies who have to report quarterly on their results, so having a team with a system in place to support that at the highest level is very important.*"

FIGURE 1

Source: NCSA and Hyperion Research, 2020

The 50-member NCSA Industry team works across domains beyond just traditional HPC modeling and simulation. One common element to these partner efforts is their need to derive value from ever-expanding volumes of data. According to McGinty, this Big Data trend along with the rise of AI on GPUs has fueled a 400 percent revenue increase for the NCSA Industry program in the past 2.5 years.

McGinty described how the team's Industrial Data Analytics Group is designed to meet this need. With expertise in data analysis and management on a massive scale, the team helps industry partners with in-memory data analytics, industrial applications of machine learning, GIS spatial data, and the convergence of artificial intelligence with high performance computing. Their current projects include pharmaceutical, agriculture, energy, Insurance risk analysis, manufacturing, and finance.

"Our alignment with academic research is important. At the University of Illinois, we're standing up a new center for digital agriculture, so we have PIs putting together a master's curriculum for that. We're also standing up a center for artificial intelligence. What that does is it create opportunities to leverage funding sources in new ways, aligning what we're doing with industrial needs while creating scholarship opportunities for students."

FIGURE 2

Source: NCSA and Hyperion Research, 2020

McGinty described the main HPC resources offered by the NCSA Industry Team:

- The Forge cluster is an industrial Dell EMC cluster designed for the convergence of HPC and AI, with Intel Skylake Xeon processors and NVIDIA V100 GPUs. Supported by the team's rapid support and advanced consulting, the Forge system is built specifically for the HPC/AI production workloads that industry is bringing to the NCSA Industry program.
- Blue Waters, one of the largest supercomputers in academia, is at the heart of the National Petascale Computing Facility. Blue Waters provides a resource for high-end, grand challenge

types of problems. The system was nearing the end of its life at NCSA but has been newly funded for GIS projects by the NSF.

- Tier-one cloud providers like AWS, Azure, and Google also work with the NCSA Industry program in cooperation with their industrial partners.

FIGURE 3

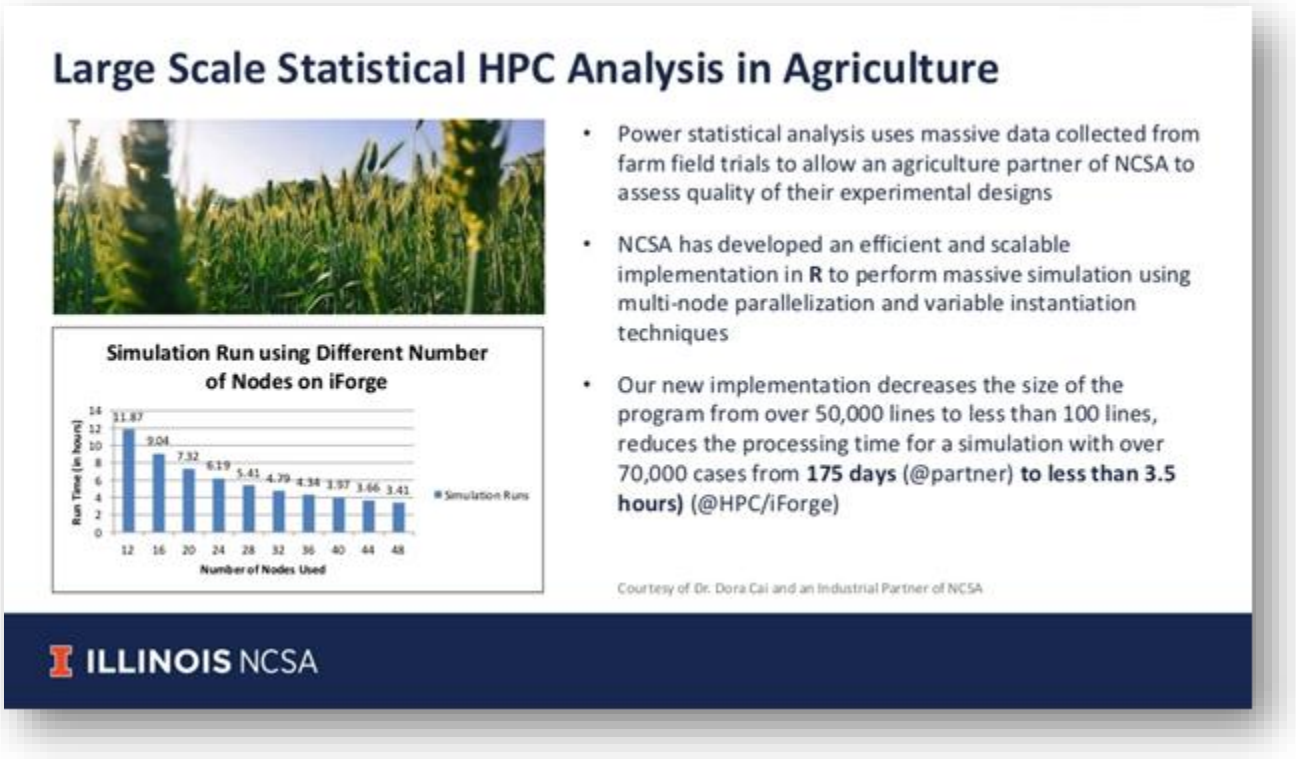

Source: NCSA and Hyperion Research, 2020

To illustrate the computational power of Forge, McGinty described an agricultural case study with an NDA agriculture customer centered around High Performance Data Analytics (HPDA).

The customer was faced with the challenge of deriving insight from massive data from farm field trials. NCSA developed an efficient and scalable implementation in R to perform a massive simulation using multi-node parallelization and variable instantiation techniques. NCSA's new implementation decreased the size of the program from over 50,000 lines to less than 100 lines and reduced the simulation runtime to just 3.5 hours.

"In this particular case, what we have is a statistical power analysis and being able to go to make code more efficient, optimize it, and take a workload that was running in 175 days on the partner's premise and reduce that down to three and a half hours on Forge. Here's what gets me excited is what does that mean? Well, to them it meant $30 million of return on investment annually."

In this proof-of-concept for the partner, NCSA was also able to reduce mean absolute errors from 0.73 to 0.43. The goal of the project was to help the partner realize 0.05 reduction in errors, so NCSA far-exceeded expectations by achieving a net 0.3 reduction.

FIGURE 4

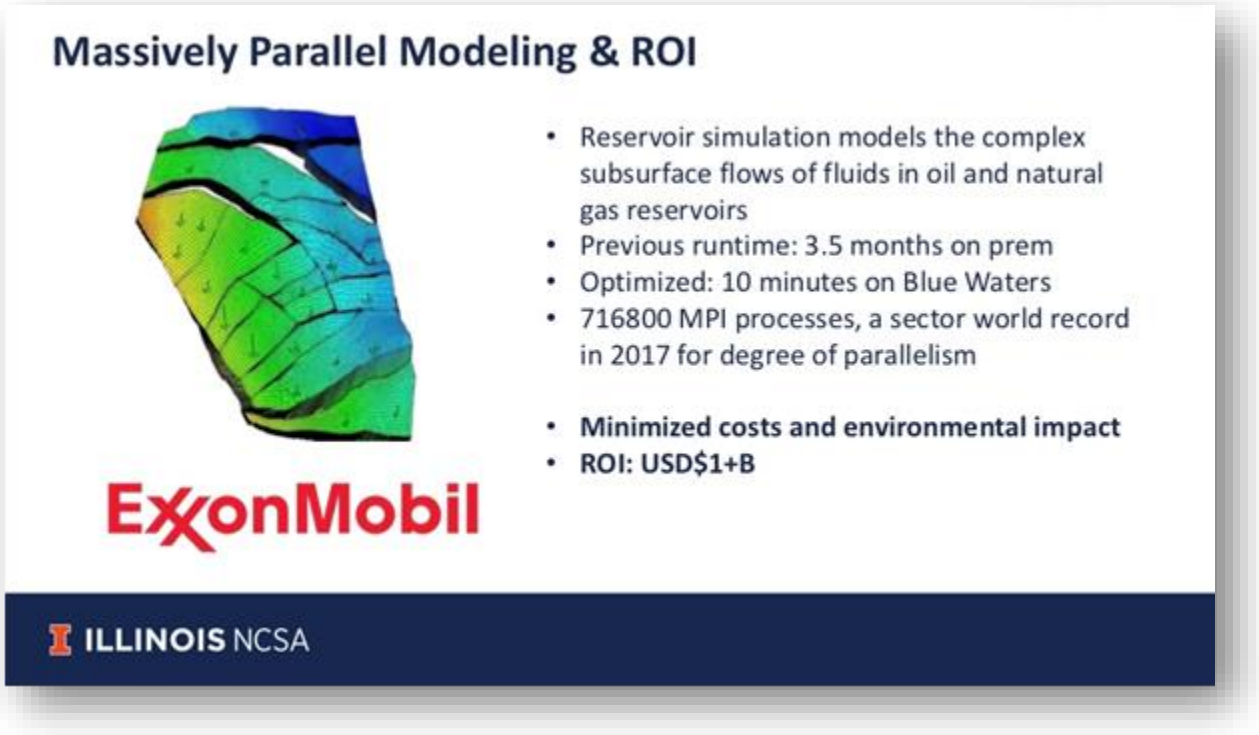

Source: NCSA and Hyperion Research, 2020

McGinty presented another case study with ExxonMobil, a simulation of complex subsurface fluid flow in oil and natural gas reservoirs. NCSA brought their experience with large scale MPI jobs to bear and was able to run the reservoir across the entire Blue Waters supercomputer. For a simulation that took 3.5 months for ExxonMobil to run on-prem in Houston, NCSA was able to reduce the run time to just 10 minutes on Blue Waters.

Over the course of this project, NCSA was able to set a 2017 record for degrees of parallelism with a total 716,800 concurrent MPI processes. With this kind of speed increase, energy companies like ExxonMobil can significantly reduce costs and environment impact.

"Now this is two years old, but it hits very broadly and what this meant as far as ROI. You see at the bottom over a billion dollar ROI as a result. We did not calculate that for them. They told us that was the result. When you go out to 716,000 MPI threads, you can go from three and a half months runtime

on-prem in Houston to running all of Blue Waters in just 10 minutes. What a result! We have an opportunity we can't take for granted as far as what this kind of ROI means to companies."

Conclusions and Issues Going Forward:

McGinty wrapped up by summarizing the elements that make the NCSA Industry team successful:

- It's about the talent first, compute second
- Be consultative – and listen first
- Provide custom solutions – no two companies are alike
- Respect their time, or lack thereof
- Leverage resources
  - University (faculty, students, research park, et al.) – Collaborators/vendors/partners
  - Other centers (NLs, international SCs)
- There is enough opportunity to go around!
- Never settle -- keep pushing, testing, being fearless

McGinty concluded with an analogy about the rapid pace of technology change and how customers lean on NCSA to be consultative with them and to tell them what's possible with all their data using AI. He then reminded the audience that NCSA is always open to new industry partnerships.

For more information or to view this and other presentations given at HPC User Forums dating back to 2008, visit www.hpcuserforum.com.

About Hyperion Research, LLC

Hyperion Research provides data driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com
