Intersect360 Research White Paper:

AI AND THE NEW HPC: DRIVING SCIENTIFIC RESEARCH IN HIGHER EDUCATION

MARKET DYNAMICS

Science Begets Science

The very nature of scientific research is that it constantly marches forward. Once one problem is solved, it unlocks the approach to a new, harder category of problem, which itself spurs investment to solve it. Consider the case of the human genome, which was first comprehensively mapped in the year 2000. This sounds like the end of a problem. "Great, we mapped it! Can we go home?" Congratulations, you've just launched the field of genomics. There is much, much more work to be done.

In the future, not only will HPC power AI, but AI will help steer HPC.

Until we reach the end of science (not likely any time soon), there will be more scientific discoveries to be made. Higher education facilities are the cornerstones of a great deal of this work. Universities and not-for-profit organizations make it part of their charter for their faculty to pursue new frontiers of human knowledge.

As a result, academia is one of the largest consumers of high performance computing (HPC) technologies and services, with over $6 billion in investment in 2018 worldwide. Applications of HPC at universities span the full range of scientific exploration, including medicine, biochemistry, energy, nanotechnology, manufacturing, and climate modeling, just to name a few. (See Figure below.)

Academic institutions house some of the largest supercomputing facilities in the world, and they are often open to agencies, businesses, and other organizations. But even beyond these world-leading systems, innovation occurs on HPC clusters of all sizes. Most universities rely on HPC at some level in some department, whether astronomy or physics or geology, and individual researchers tend to be the ones at the forefront of their own particular specialties. As such, HPC helps higher education advance scientific discovery at all levels.

AI and the New HPC

The unending reinvention of scientific discovery demands new tools to enable it. New HPC applications, architectures, and approaches are driven by the perpetual need to solve the next big problem. Thanks to advancements in the availability of data, the scale of HPC, and investment by hyperscale web companies, there has been a revolution in using artificial

White paper sponsored by Dell Technologies and Intel Corp. Full neutrality statement at https://www.intersect360.com/features/neutrality-statement.
P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com

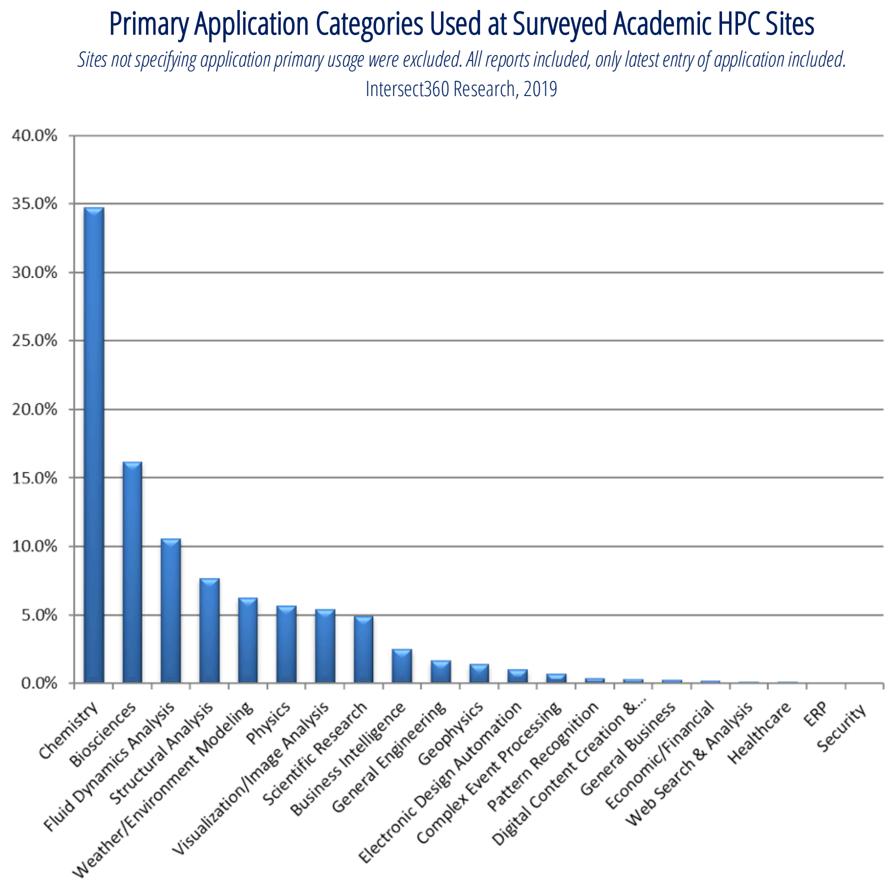

intelligence (AI)—or more specifically, machine learning—as a complement or component of HPC.

Historically, most scientific applications have been deterministic; based on a set of inputs, a program runs a bunch of calculations and comes up with an output, which is the answer. A second type of applications is probabilistic. Rather than computing one answer, we test a multitude of scenarios to see what happens. This would be the approach to answering the question, "If you shuffled together ten decks of cards at random, what is the probability that some sequence of 20 cards would contain at least 18 spades?" It is a very difficult math problem to compute exactly, but relatively simple to have a computer try it 10 million times to come up with a likely answer. In HPC, this is called a Monte Carlo simulation, and it is a common approach in some categories of scientific research.

Machine learning represents a third category of application that is experiential. Based on patterns seen previously, a machine learning algorithm makes inferences about current or future situations. This approach is called "artificial intelligence" because it mimics how humans learn. Although I have never seen this cat before, I am confident it is a cat based on my lifetime of experience and learning involving cats.

If AI can be taught the rules of games like chess or go, and within those rules, find new strategies, then can AI be taught the rules of a different "game," such as finding genetic markers of diseases, or engineering a microbe that will bind to carbon dioxide in the atmosphere, or designing polymers that are stronger and lighter?

Machine learning can be deployed any time there is a wealth of data to draw on, coupled with a reward from making more intelligent inferences based on that data. Many scientific research areas in higher education fit this pattern. HPC using organizations in general are trending toward the incorporation of AI. In a 2018 survey, 61% of HPC users said they are already running machine learning as part of their environments.

Higher education is one of the key markets pushing AI forward in science. While this can happen independently from traditional HPC approaches, more commonly it is in conjunction with deterministic scientific applications. Furthermore, HPC is in itself a key driver for AI, as certain aspects of machine learning lean on high-performance architectures.

Machine learning can be divided into two primary phases—training and inference. In the training phase, a data repository—such as sets of medical images, cosmological data, or weather patterns from past storms—are fed into a neural network, which finds patterns in the data. This phase typically requires HPC resources, including advanced processing elements. Once trained to identify the tumors, the supernovas, or the tornados, it can be turned to the inference phase, in which is makes predictions in new data sets it is given. While this approach does not replace the oncologist, astronomer, or meteorologist, it can be a vital tool in helping the expert be more focused in analyzing more data, more efficiently.

In the future, not only will HPC power AI, but AI will help steer HPC. Consider: If AI can be taught the rules of games like chess or go, and within those rules, find new strategies, then can AI be taught the rules of a different "game," such as finding genetic markers of diseases, or engineering a microbe that will bind to carbon dioxide in the atmosphere, or designing polymers that are stronger and lighter? This incorporation of AI and analytics will define the next generation of discovery with the new HPC.

INTERSECT360 RESEARCH ANALYSIS

Today's challenges in higher education are representative of the "new HPC," combining scientific computing with large-scale data analytics and artificial intelligence. Successful solution providers will be those that have technologies that can span this expanded workflow, combined with domain-specific expertise in helping organizations achieve meaningful breakthroughs.

Dell Technologies is just such a company. With trusted products across both computation and data management, Dell Technologies is the industry leader in total HPC solution revenue. Dell Technologies leverages this breadth of offerings with converged solutions that incorporate HPC, data analytics, and AI[1] and offers tailored solutions incorporating the latest in AI for a wide range of scientific domains.[2]

[1] https://www.dellemc.com/en-us/collaterals/unauth/brochures/solutions/hpc_ai_convergence_brochure.pdf.

[2] https://www.emc.com/collateral/solution-overview/ready-solns-for-ai-machinedeep-learning-sol-overview.pdf.

Most importantly, Dell Technologies is helping higher education derive real value from their investments in HPC, analytics, and AI. Purdue University uses a "community cluster" model responds to ever-evolving demands from its scientific community, operating as many as five highly scalable supercomputers in any given year.[3] Two of these are currently supplied by Dell Technologies, including the "Brown" clusters with over 550 Dell server nodes. Preston Smith, Purdue's director of research computing services and support, describes the challenges of meeting the needs of over 1,300 individual researchers:

> *As our organization builds systems for artificial intelligence and machine learning workloads, we are seeing researchers coming from domains that we have not previously served. For example, people solving problems in humanities with text analysis, or people doing image processing — they need larger systems with more storage and larger amounts of data to get results faster.*

Dell Technologies also powers the "Pitzer" system at Ohio Supercomputer Center (OSC)[4] and the "Great Lakes" cluster at the University of Michigan.[5] Both of these handle diverse and emerging workloads in HPC and machine learning. At Michigan, about half of the research is directed toward health, in areas such as personalized medicine and genetics, while the other half is spread across advanced manufacturing, smart cities, autonomous driving, and other scientific research. OSC's Pitzer is equally diverse, supporting HPC and machine learning research across agriculture, precision medicine, structural mechanics, and natural sciences.

"OSC went with Dell EMC for its last two major clusters because of the variety of technology partners that they can bring to the table, the integrated solutions they can deliver, the price of the system and the support they provide," says OSC Executive Director David Hudak. "They provide excellent value and terrific support."

At the high end, Dell Technologies powers much of the supercomputing infrastructure at the Texas Advanced Computing Center (TACC) at the University of Texas at Austin, including the newest "Frontera" supercomputer. Frontera includes the latest Intel Xeon Scalable processors for HPC to reach over a petaflop of computing performance. Funded by NSF, "Up to 80% of the available hours on Frontera, more than 55 million node hours each year, with be made available through the NSF Petascale Computing Resource Allocation program."[6] Dell is also behind the Stampede2[7] and Wrangler[8] supercomputers at TACC, with applications spanning scientific computing, business computing, and big data and analytics.[9] Across these platforms,

Research scientists at the Texas Advanced Computing Center (TACC) leverage the capabilities of second-generation Intel Xeon Scalable processors to enable new kinds of scientific discoveries and engineering research.

[3] Dell Technologies customer case study, 2018.
[4] Dell Technologies customer case study, 2018.
[5] Dell Technologies customer case study, 2018.
[6] https://www.tacc.utexas.edu/systems/frontera.
[7] https://www.tacc.utexas.edu/systems/stampede2.
[8] https://www.tacc.utexas.edu/systems/wrangler.
[9] https://www.emc.com/collateral/customerprofile/2017-tacc-10023036-wrangler-hpc-data-analytics.pdf.

research scientists at the TACC leverage the capabilities of second-generation Intel Xeon processors to enable new kinds of scientific discoveries and engineering research.

For higher education, the new HPC is fueled by analytics and AI. Across scientific domains, there is a deep wealth of data. If harnessed, it can unlock new discoveries, which themselves will spawn new areas of research. While the fundamental drivers of research remain unchanged, these converged, high-performance solutions are critical to enabling new generations of insight. With its domain-specific knowledge and technology solutions across data management and computation for HPC, analytics, and AI, Dell Technologies is well-positioned to help is customers in higher education achieve new insights for the next generation.

To learn more, please visit DellEMC.com/HiEd.

P.O. Box 60296 | Sunnyvale, CA 94088 | Tel. [888] 256-0124
www.Intersect360.com | info@Intersect360.com