

Alex Wright

Big Data Meets Big Science

Next-generation scientific instruments are forcing researchers to question the limits of massively parallel computing.

N A SECLUDED hilltop outside Palo Alto, CA, Jacek Becla leads a team of researchers at the SLAC National Accelerator Laboratory who are quietly building one of the world's largest databases.

Scheduled to go live in 2020, the Large Synoptic Survey Telescope (LSST) will feature a 3.2-gigapixel camera capturing ultra-high-resolution images of the sky every 15 seconds, every night, for at least 10 years. Ultimately, the system will store more than 100 petabytes (about 20 million DVDs' worth) of data, but that is barely a fraction of the data that will actually pass through the camera.

"Even though we are dealing with huge amounts of data, there is even more data that we are not saving," says Becla. With 40 billion–50 billion potential astronomical objects in the camera's purview, he explains, it would be all but impossible to store every pixel in perpetuity. Instead, the system will extract critical data from the images in real time, then simply discard the source images.

As increasingly powerful large-scale scientific instruments come online from the Large Hadron Collider to advanced light beam processors and mo-



Argonne National Laboratory chemist Karena Chapman peers inside the vacuum tank of the new high-energy Si Laue monochromator recently installed in the Argonne Advanced Photon Source, an upgrade that increased the X-ray flux (the number of photons focused on the sample being studied) by a factor of 17.

lecular imaging tools—they are starting to churn out more data than even the most powerful massively parallel supercomputers can handle. As a result, scientists are exploring new approaches to reducing those datasets to manageable size, incorporating new learning from the private sector about cloud-based computing, and in a few cases exploring the possibilities of emerging frameworks like quantum computing.

Those strategies stand in stark contrast to the traditional scientific approach to high-performance computing, which has long relied on a "brute force" approach involving stringing together greater and greater numbers of CPUs and disk arrays. After a decadeslong infatuation with parallel supercomputing, however, some researchers are beginning to butt up against the limitations of that approach.

"Moore's Law is effectively already broken down," says Massachussetts Institute of Technology professor Scott Aaronson, who argues the laws of physics have caught up with Intel founder Gordon Moore's famous dictum that the number of transistors on integrated circuits would double every two years.

Researchers also are grappling with both economic and algorithmic constraints that force them to explore methods beyond the tried-and-true technique of throwing ever more processors at a problem.

At the Argonne National Laboratory in Illinois, Chris Jacobsen leads a team working on the Advanced Photon Source (APS), an enormous, football field-sized synchrotron that produces X-ray photons by swirling electrons around a circular apparatus at nearly the speed of light. Researchers from 65 different field stations rely on the machine to gather imaging data about a wide range of subjects: from proteins and nanoribbons to lithium-ion batteries and catalytic converters.

The experiments vary tremendously in scope, but the data they collect always comes in intense bursts of up to 11 gigabytes of raw data per minute. In a typical month, APS distributes about 112 terabytes of data. "We get so much data that we can't just sit there and examine it by hand," says Jacobsen. "It takes time for all these processors to communicate with each other, and they can't send messages faster than the speed of

"For most applications, the real bottleneck is not the processing time," but "the need to constantly retrieve stuff from memory."

light-that's sort of the ultimate limit."

Given those constraints, the team is constantly looking for more efficient ways to help researchers interpret their test results. "What are the features we can pull out from the raw data? What can we understand, rather than just measure?"

With so many far-flung researchers, Jacobsen's team has also been wrestling with how best to deliver datasets to its many end users. In the past, several teams relied on their own ad hoc "sneakernets," lugging their hard drives to the facility for a few days before bringing them back home. As the lab continues to improve its detectors, however, the data rates keep going up, forcing the Argonne team to explore new cloud-based approaches to providing data to researchers.

"We are trying to move towards a more cohesive computing strategy," says Jacobsen.

Recently, Argonne's physicists have been collaborating with their colleagues from the applied math and computer science departments to develop new tools to allow researchers to automate the transfer of data from the beamline computer to a central data store where it can be optimized, backed up, and managed. That data is then made available via a secure TCP/ IP connection, using a tool called Globus Online (globus.org), and stored using Amazon Web Services—allowing for multiple parallel connections.

In a similar vein, researchers at the U.S. Department of Energy's Brookhaven National Laboratory are exploring cloud-based approaches to harnessing the vast troves of data currently being produced by the ATLAS experiment at Europe's Large Hadron Collider (LHC), famous for discovering the elusive Higgs boson.

ATLAS has already generated 140 petabytes of data, distributed between 100 different computing centers, with most of it concentrated in 10 large computing centers like CERN and Brookhaven.

Physicist Alexei Klimentov has been working on a framework for managing this enormously complex computational enterprise—which involves an estimated 3,000 physicists creating more than two million computing jobs per day—using a system called PanDA (Production and Distributed Analysis),

"PanDA is a pilot system," says Klimentov. "It knows about the site, the software, the storage, and the available CPU slots. Then, according to the available resources, it matches them against the payload." For example, a simulation project typically requires a lot of processing but little data storage, whereas a complex data analysis job requires fast access to large hard drives.

By distributing these jobs across the cloud to the most appropriate available system, PanDA can make the best use of available resources while minimizing system downtime. Even so, shuttling 140 petabytes of data around the world is no small undertaking.

Engineering fast and reliable data transfer mechanisms is emerging as one of the critical challenges for scientists working with big data—not just for moving files from computer to computer, but for shuttling data in and out of memory as well.

"In traditional high-performance computing, you have very little data and very little I/O," says Becla, "so you are basically reading the data into memory and doing the processing in memory. But in the big data world, you cannot do this. You cannot have a trillion pieces of data in memory at the same time."

MIT's Aaronson echoes that concern. "For most applications, the real bottleneck is not the processing time," he says, but "the need to constantly retrieve stuff from memory. For a lot of programs, the processor is sitting idle, waiting for the memory to come back." The challenge, then, is how to design memories that are fast, large, and responsive.

Aaronson likens the problem of classical computing to the eternal conundrum of finding an apartment in New York City: "You could get it in Manhattan where it's close to everything but small and expensive, or you go to Long Island and it's cheaper but farther from everything." Similarly, computer scientists must navigate complex trade-offs in trying to optimize system performance with large datasets. "Registries are super-scarce, then you go out to L1 and L2 cache and then the RAM, then you're in the boonies of the hard disk. How do you optimize the trade-off?"

That tension captures the challenge of cloud computing: how to take advantage of the economies of the cloud without losing the gains of having everything in close proximity?

"You can just throw more parallelism at things, but the amount of memory and the amount of disk space has been blowing up tremendously," Aaronson says. "Even if in principle you have all these parallel processes, it can be harder to write code that takes advantage of the parallelism."

The traditional approach to highperformance computing relied on millions of CPUs to perform many calculations on relatively small chunks of data. Up until recently, most largescale systems fell into this category. Yet in the scientific world, where data is increasingly interrelated, problems are becoming tougher to parallelize.

Some researchers hold out hope for quantum computing, a much-hyped field that promises enormous computational speed gains. However, Aaronson advises caution. "There's a temptation for people to look at quantum computing and say, 'this must be the thing that will continue Moore's Law,' but a lot of that relies on misconceptions about what a quantum computer is. It's a fundamentally different kind of computer."

Unlike a classical computer that can perform a large number of calculations at the same time, quantum computers rely on subtle effects from quantum mechanics that can solve certain classes of problems much faster; for example, breaking cryptographic codes, factoring large numbers, or simulating quantum physics. For more traditional computing tasks, like combinatorial optimization, airline scheduling, or adiabatic algorithms, it is not at all clear that quantum computers will offer any meaningful performance gain.

"It's conceivable that a quantum computer could help with protein fold-

ing or DNA sequencing, but the advantages are not obvious," says Aaronson. "You'll get an advantage from a quantum computer only when you can figure out how to exploit quantum interference."

In the near term, scientific researchers may take solace in knowing they are scarcely alone in grappling with the challenges of big data. The explosive growth of the consumer Internet has thrust many of the Internet's leading companies into similar territory.

In 2007, Becla organized a 60-person workshop called the Extremely Large Database group (XLDB), which has since grown into a network of more than 1,000 members spanning numerous scientific research centers, as well as private-sector participants from Google, Amazon, eBay, LinkedIn, Yahoo!, and elsewhere.

Increasingly, these organizations find themselves operating in overlapping territories: working with large collections of images, time series, or determining how best to detect outliers in large datasets, whether in the form of gamma-ray bursts or security intrusions.

"We see commonalities between what astronomers are doing and what eBay and Wall Street are doing," says Becla.

Who would have thought that the path to unlocking the mysteries of the universe might run through eBay? Says Becla, "It's an eye-opener."

Further Reading

Francesco De Carlo Dóga Gürsoy et al. Scientific Data Exchange: A schema for HDF5-based storage of raw and analyzed data. Submitted to J. Synchrotron Radiation.

Rachel Mak, Mirna Lerotic, Holger Fleckenstein, Stefan Vogt, Stefan M. Wild, Sven Leyffer, Yefim Sheynkin, and Chris Jacobsen. Non-negative matrix analysis for effective feature extraction in X-ray spectromicroscopy. Submitted to the Royal Society of Chemistry for a Faraday Discussion meeting. DOI: 10.1039/ c000000x

Michael Stonebraker, Paul Brown, Donghui Zhang and Jacek Becla.

SciDB: A database management system for applications with complex analytics. *Computing in Science & Engineering*, 15, 54-62 (2013), DOI:http://dx.doi.org/10.1109/ MCSE.2013.19

 $\ensuremath{\textbf{Alex}}\xspace \ensuremath{\textbf{Wright}}\xspace$ is a writer and information architect based in Brooklyn, NY.

© 2014 ACM 0001-0782/14/07 \$15.00

ACM Member News

SECURITY NEEDS TO BE FLUID, PRAGMATIC: YUNG



"Cryptography theory is great," says Moti Yung, a Google research scientist and 2013 ACM

Fellow. "There are terrific security tools and they work well, assuming no one has penetrated your computing systems." In an era of global and mobile systems, however, Yung believes businesses need more pragmatic and fresher approaches to secure systems already "partially penetrated."

Yung knows what he is talking about; he has over 30 years' experience as a cryptographer. He earned his Ph.D. from Columbia University, where he currently serves as an adjunct professor and visiting senior research scientist. Prior to joining Google in 2007, he was a cryptographer and scientist at IBM's Thomas J. Watson Research Center; chief scientist at CertCo, and director of Advanced Authentication Research at RSA Laboratories.

In 2004, Yung and Adam Young co-authored "Malicious Cryptography: Exposing Cryptovirology," a book focusing on cryptovirology (the use of rogue code to invade systems) and kleptography (a targeted stealth-attack mechanism using a crypto system to try to break or defeat another cryptographic system from within).

Privacy in the cloud computing era brings new challenges daily, Yung says. Ongoing and growing security threats make privacy and securing data crucial.

"It makes no sense to build a Maginot line that people can circumvent; security systems need to be fluid," Yung says. "Security professionals must thrive on challenges, accept failures as learning experiences, and move on." That is why he monitors hackers and continually updates his security research directions. His most significant accomplishment, Yung says, is "still in the future. You can't rest on your laurels."

–Laura DiDio