Update

Disruptive Technologies in High-Performance Computing: HPC User Forum, April 7–9, 2014, Santa Fe, New Mexico

Earl C. Joseph, Ph.D.
Steve Conway
Chirag Dekate, Ph.D.

IN THIS UPDATE

This IDC update captures part of the proceedings at the 52nd HPC User Forum held April 7–9, 2014, in Santa Fe, New Mexico. Earl Joseph, HPC program vice president, IDC, moderated a disruptive technologies panel. This update is a near-verbatim record of the panel.

Disruptive Technologies Panel — Moderator: Earl Joseph, IDC

- **Bo Ewald, D-Wave:** D-Wave was founded in 1999. We do zero flops. We have 5 bits of precision. Our quantum computer runs at almost absolute zero. It uses only 15.5kW. We have two 512-qubit systems, one installed at Lockheed (housed at USC) and one installed at NASA Ames (for Google). We have demonstrated performance 10,000–100,000 times faster than classical computers. Our computer natively solves an optimization problem involving machine learning and Monte Carlo/sampling (e.g., for FSI). The computer uses quantum mechanical effects, quantum algorithms, and qubits rather than bits.

  It's cooled to 0.02K, 150 times colder than interstellar space, and is shielded to 50,000 times less than Earth's magnetic field. We use standard fabrication processes. There is just one machine instruction. You provide the values and tell it how many solutions you want. A typical number would be 10,000 solutions in a second. Google is really a machine-learning company. It came up with an algorithm for our machine to recognize cars with 94% accuracy. Another problem Google gave us to run was discrete combinatorial optimization benchmarks. Only with the D-Wave, the solution time remained low no matter how large the problem became. We're trying to double the number of qubits every year. We believe the D-Wave system will work side by side with classical HPC systems and data analytics systems.

- **Rishi Khan, Extreme Scale Solutions:** We are software developers for big compute and big data. On the research side, we build software stacks for next-generation government HPC systems. We also do proof-of-concept projects for pioneering initiatives. We are building a software product called DiscoveryMiner for turning data into insights. This is aimed at the average enterprise. For example, look at the 2013 Major League Baseball statistics. Our solution will tell you things like, if the game goes on really long, the away team has a substantial advantage. We are looking for use cases from the HPC community (i.e., data problems you want solved).

- **Bob Keller, Silicon Informatics:** We develop parallel pseudorandom number generators for large parallel computations. Modelers understand the dangers of coherence in random number generators. We've developed a correlation tester to evaluate the numbers from any

random number generator across parallel streams – up to 1.5 billion streams so far. We're looking for collaborations to evaluate our generator and tester. We're looking for DoD and DOE users.

- **Daniel Hardman, Adaptive Computing:** You probably know us as the makers of the Moab scheduler. We made some strategic investments, and we are changing the nature of scheduling. Centralized schedulers that make all the decisions have reached their limits. Our new MTM product lets you take a certain category of work and, say, you schedule that and let the main scheduler schedule the rest of the work. We also have been working on data staging. We think of the future in terms of "big workflow" that may be enterprise, HPC, cloud, big data, or other things.
- **Bill Mannell, SGI ICE X Immersion System:** The technology is phase-change cooling technology. Development was funded by DoD. We worked with Berkeley Lab, 3M, Intel, NRL, and others. An SGI ICE X supercomputer is working today with this in a 3M lab. There's no need for cold plates or a closed-loop cooling system. No changes are needed to a standard air-cooled computer. You immerse the computer in this liquid and the hottest components trigger a phase change that cools them.
- **Bill Nitzberg, Altair:**We do PBS Works. We're doing a lot with power, scalability, and resilience. In scale, we are in the camp that says parallelism is the key. In resilience, we built an entire framework to support health monitoring of systems. In power, at NASA, we made slides on a beautiful color printer and a not-so-beautiful printer. I switched to the non-beautiful printer when NASA started charging money to use the beautiful printer. If we are given information on the fly about our power use in HPC, I think we will use less.
- **Leo Reiter, Nimbix:** We focus on HPC as a cloud service. We serve industries that have massive data problems. We are a data-processing cloud, not a usual cloud. Today, you can use an on-premise cluster or you can access a commodity cloud that's designed for Web services, but not for HPC. There are brokers who can adapt your problem to these clouds, and that adds costs. Ours is the Nimbix Accelerated Compute Cloud. Users submit large processing tasks via a Web application, from any device. Workloads include bioinformatics, transcoding, simulations, and rendering. Our cloud has Xeon Phi, FPGA, GPGPU, DSP, and CPU.
- **Dale Southard, NVIDIA:** We're introducing stacked memory for much higher bandwidth (about 1TBps) and lower cost per bit. We're also introducing NVLink, a GPU high-speed interconnect with speeds of 80–200GBps and planned support for Power CPUs. NVLink enables data transfer at the speed of CPU memory. HPC users can ride the server market, the embedded market, or the GPU market.
- **Mike Vildibill, DDN:** Our disruptive technology is HPC burst buffers. They require smaller, highly robust parallel file systems with very high bandwidth efficiency. In tests at Argonne, 99% of the time, storage bandwidth utilization was less than 33% of max, and 70% of the time, it was less than 5% of the max. The burst buffer absorbs the peak load, while the file system handles the remaining load. Burst buffer is a layer of SSDs. They use I/O acceleration, out-of-core I/O, data analysis support, and temporary data storage.
- **Leijun Hu, Inspur**:A group in our company is doing something interesting. The human brain has its own way of computing that uses fuzzy logic. We must make future semiconductors perfect at increasingly challenging nm dimensions. We have the opportunity to design a silicon brain that operates more like the human brain. This would involve machine learning.

- **Barry Bolding, Cray:** From a Cray perspective, a big question involves the wider consequences of single disruptive technologies. Also, what's most needed to bring the technologies to market faster or with more certainty?

Question 1: What Do You Most Need to Bring These Disruptive Technologies to Market Faster or with More Certainty?

- **Extreme Scale:** Cash and partnering with the right people.
- **Nimbix:** ISVs, especially ones that will embrace the cloud with appropriate licensing.
- **Cray:** It certainly takes understanding of the business challenges. "Coolness" is not enough.
- **Altair:** I'd echo that.
- **SGI:** For 2-phase cooling, we need people who are worried about their cooling challenges.
- **DDN:** Confidence is important, so test beds and other collaborations make things go faster.
- **IBM:** People need to understand that the paradigms they grew up with may no longer apply.
- **Silicon Informatics:** We need cash, user validation, and partnerships, along with other applications that need our technology.
- **Adaptive:** I hope OpenStack will contribute to the Big Workflow technology in HPC.
- **NVIDIA:** We need customers with problems.
- **Inspur:** We need more people to get involved in exploring our disruptive idea.
- **D-Wave:** We need a handful of users who want to be pioneers in the non-von Neumann world and also more cash to make our system easier to use.

Question 2: Are There Any Particular Partners You Are Seeking, and Can the HPC User Forum Help Develop These Partnerships?

- **Inspur:** Most HPC innovation has come from the United States. The HPC User Forum is helping us to meet a wider community, including potential partners.
- **NVIDIA:** We partner with OEMs and customers to meet their needs.
- **Adaptive:** We're looking for people interested in better performance in exchange for less precision, for exploratory purposes.
- **Silicon Informatics:** The HPC User Forum has already helped us in our R&D.
- **IBM:** For the microserver, we'd like an SoC partner. The HPC User Forum has already helped today. We want partners open to getting away from classic models for problem solving.
- **DDN:** There are endless opportunities for partnering around SSDs and future I/O architectures.
- **SGI:** For phase-change cooling, we want more users willing to be partners to help us advance the technology to the next level.
- **Altair:** The HPC User Forum has helped us develop more partnerships. We need people who can tell us more about the problems they need to solve and give us workload data related to this.
- **Cray:** The User Forum works pretty well. We could get smaller groups together in the User Forum. Regarding direct partnerships, we do need data to see how things are working when they are not working.

- **Nimbix:** We want ISV partners interested in the missing middle problem.
- **Extreme Scale:** I think this User Forum contains a lot of opportunities for partners.

"Storage Requirements for the Future: It's the Applications That Matter": Henry Newman, Instrumental

I want to talk about the weather workflow example and its impact on the storage stack and how we need to rethink data analysis. Today, weather data is collected by satellite. There are raw images and multispectral images, and you also have fixed buoys, plus land observations and radar around the world and collections from planes and ships and more. Tomorrow, data will come from home weather stations, cell phones, new satellites with better resolution, new higher-resolution radars, etc. So there's more data coming in, and it's not big, bursty data. It's all small stuff, with 32-bit values. Data's collected for each grid point and for multiple attributes: pressure, temperature, moisture, wind velocity, and direction. This is all put into a data structure, by grid point.

These folks often don't know about alignment on file systems, RAID, or disk boundary. Therefore, I/O is used inefficiently. Some of these data structures are decades old and in some cases it's harder to redo them than to update the models. This all leads to a major data assimilation problem, where you could even be running your weather forecast against old data. Even with the best weather model, if you can't assimilate the new data fast enough, you may be using outdated initial conditions and produce a forecast that's off.

Once you assimilate the data, you then have to run it into the model. There are various ways to do this. Requirements:

- Fast metadata transaction rates (>50,000 transactions per second)
- High aggregate rates from many independent parallel threads
- Directory-based quotas that help some environments
- Delivery of a large fraction of file system performance metrics when disks are rebuilding

Future methods:

- 4D-Var is for observations that are distributed in time.
- Throwing more hardware at unaligned data will not solve the problem.
- The POSIX stack is not going to change.
- Small-block I/O does not do well on any file system.
- Hardware alignment issues are a big problem for storage access.

There are many other data problems that involve small blocks. What has to happen: The only solution I see is a complete restructuring. Today, people are not thinking much about restructuring data. People will need to start dealing with REST interfaces because that's where the market is going. Predicting the impact of rainfall and temperature on future crops is worth trillions of dollars.

LEARN MORE

Related Research

Additional research from IDC in the technical computing hardware program includes the following documents:

- *Worldwide High-Performance Data Analysis 2013–2017 Forecast* (IDC #241315, June 2013)
- *Experiences in HPC: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241455, June 2013)
- *HPC in Aerospace, Astrophysics, and Astronomy: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241451, June 2013)
- *Top Issues for HPC Sites: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241463, June 2013)
- *Potential Disruptive Technologies Panel: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241452, June 2013)
- *Advanced Visualization: HPC User Forum, April 29–May 1, 2013, Tucson, Arizona* (IDC #241446, June 2013)
- *Worldwide Technical Computing Server 2013–2017 Forecast* (IDC #241154, May 2013)
- *Supercomputers Exceed 50% of the HPC Server Market in 2012* (IDC #240426, April 2013)
- *Changing Market Dynamics: HPC Meeting Big Data and IDC's Projected Evolution of the Market* (IDC #240365, March 2013)
- *Livermore Lab Expands Industry Partnerships: Economic Security Is Vital for National Security* (IDC #240232, March 2013)
- *Worldwide Technical Computing 2013 Top 10 Predictions* (IDC #239421, February 2013)
- *High-Performance Data Analysis at NASA JPL* (IDC #238254, December 2012)
- *Advanced Research at the South Ural State University Supercomputing Center* (IDC #238225, December 2012)
- *The Economic Value of HPC in Science and Industry: HPC User Forum, September 2012, Dearborn, Michigan* (IDC #237182, October 2012)
- *Big Data in HPC: HPC User Forum, September 2012, Dearborn, Michigan* (IDC #237180, October 2012)
- *How Nations Are Applying High-End Petascale Supercomputers for Innovation and Economic Advancement in 2012* (IDC #236341, August 2012)
- *Tokyo Institute of Technology: Global Scientific Information and Computing Center* (IDC #236243, August 2012)
- *Shanghai Supercomputer Center* (IDC #236245, August 2012)
- *Petascale Supercomputing at the University of Tokyo* (IDC #236292, August 2012)

- *HPC Application Leadership at the Supercomputing Center of Chinese Academy of Sciences* (IDC #236281, August 2012)
- *India Broadening Access to Supercomputing* (IDC #IcUS23621912, July 2012)
- *Institute of Process Engineering, Chinese Academy of Sciences* (IDC #236204, July 2012)
- *How Fujitsu Built the World's Fastest Supercomputer in Record Time and Ahead of Schedule* (IDC #235733, July 2012)
- *Europe Sharpens Its Focus on Exascale Computing* (IDC #IcUS23555112, June 2012)
- *IBM Returns Supercomputer Crown to United States* (IDC #IcUS23547812, June 2012)
- *Potential Disruptive Technologies: HPC User Forum, April 2012, Richmond, Virginia* (IDC #234742, May 2012)
- *The Broader HPC Market: Servers, Storage, Software, Middleware, and Services* (IDC #234682, May 2012)
- *HPC End-User Site Update: RIKEN Advanced Institute for Computational Science* (IDC #233690, March 2012)
- *National Supercomputing Center in Tianjin* (IDC #233971, March 2012)

- *HPC End-User Site Update: RIKEN Advanced Institute for Computational Science* (IDC #233690, March 2012)
- *National Supercomputing Center in Tianjin* (IDC #233971, March 2012)

- *Worldwide Data Intensive-Focused HPC Server Systems 2011-2015 Forecast* (IDC #232572, February 2012)

About IDC

International Data Corporation (IDC) is the premier global provider of market intelligence, advisory services, and events for the information technology, telecommunications and consumer technology markets. IDC helps IT professionals, business executives, and the investment community make fact-based decisions on technology purchases and business strategy. More than 1,100 IDC analysts provide global, regional, and local expertise on technology and industry opportunities and trends in over 110 countries worldwide. For 50 years, IDC has provided strategic insights to help our clients achieve their key business objectives. IDC is a subsidiary of IDG, the world's leading technology media, research, and events company.

Global Headquarters

5 Speen Street
Framingham, MA 01701
USA
508.872.8200
Twitter: @IDC
idc-insights-community.com
www.idc.com