

November 07, 2014 | By Doug Ramsey

UC San Diego Computer Scientists Triumph at Data Processing Competition



Computer science and engineering professor George Porter

After a two-year hiatus, a team from the Center for Networked Systems (CNS) at the University of California, San Diego came roaring back to set three new world records in a data processing competition for industry and academe. CNS associate director George Porter, former CNS director Amin Vahdat (now at Google), and Computer Science and Engineering Ph.D. student Michael Conley set a world record in the 100 Terabyte Daytona (think speed) GraySort category. They outperformed everyone else, sorting 100 TB in

less than 23 minutes, but tied with the startup company Databricks (which sorted the same amount of data in 23.4 minutes). Both used the Amazon Elastic Compute Cloud (EC2).

“They used 10 percent more processors to achieve very close to the speed that we did,” said Porter. “But under the rules, the speed must be substantially faster, so we tied for first place and we’re fine with that.”

This week organizers of the Sort Benchmark competition also announced that the UCSD team clinched two other world records in the first-ever 100 TB CloudSort competition. Entries were required to do all of their data processing on publicly-available cloud services instead of on dedicated computer clusters. Porter and Conley were the unequivocal winners in the two new cloud categories.

“The results underline our growing emphasis on creating experimental artifacts that advance the state of the art in practice based on research done at UCSD,” observed Rajesh Gupta, chair of the CNS department in the Jacobs School of Engineering.

For Porter and Conley, the wins were particularly nostalgic because they and Vahdat were part of the original CNS team that clinched multiple back-to-back world records in the data-sorting challenge in 2010 and 2011, working with then-graduate student Alex Rasmussen (Ph.D. '13). Conley and Rasmussen developed what was, at the time, the world's fastest sorting system, called TritonSort, which achieved record speeds by focusing on per-disk and per-node efficiency.



CSE Ph.D. student Mike Conley

Ironically, in the 2014 competition, China's giant web services company Baidu was the winner in the 100 TB Indy GraySort category, but as Porter noted, "the plot thickens ... they ran TritonSort!" Baidu scientists reimplemented the system developed in 2011 at UCSD, but they renamed it BaiduSort and ran it on a cluster of nearly 1,000 servers (compared to the 52 machines in the CNS cluster in Atkinson Hall that produced the 2011 record).

CNS didn't enter the competition in 2012 or 2013 because their dedicated computer cluster was no longer sufficient to win records. "We had maxed out our ability to continue setting records unless we were willing and able to spend millions of dollars to expand the cluster itself," recalled Porter. "We were up against very well-heeled competitors, so throwing money at the problem was not going to work."

While Conley had been working on big-data processing since 2009, he and Porter revved up for this year's competition starting last January. Their goal: to transition from data sorting using the dedicated cluster (which yielded their earlier record speeds) to the new world of cloud-based data sorting.

When organizers of the data-sorting competition announced the new cloud-based categories, the CNS team – and Conley in particular – focused on finding ways to be better than the competition. And not just against their rivals. "We wanted to prove that we could do as well with cloud resources as we had been able to do when we set a world record using the dedicated cluster," said Conley. "In the end, our speed was nearly ten times faster, and we were able to almost double the efficiency from our 2011 record. In my opinion, that is our big victory – improving our efficiency compared to our own previous record."

Topping the previous record also reinforced the dramatic new cost paradigm in cloud computing. In the final competition, the UC San Diego team calculated that the winning entry used only \$451 worth of Amazon EC2 computing time, in contrast to the dedicated, multi-million-dollar cluster the team used to set their 2011 records.

“What’s changed and what’s different is that you are starting to see a lot of companies, large and small, research organizations and the federal government are moving to cloud computing because there is no capital expense, you just have to pay for what you use, and that’s a very different set of assumptions than what we’re dealing with,” said Porter, adding that it’s different in two fundamental ways. “The first is that we don’t control it anymore. You’re on the highway with all of these people and some are good drivers and some are not-so-good drivers, and there are potholes and lots of other problems. The other challenge was more subtle: the actual functionality that the cloud providers give you is a function of the least common denominator functionality that their customers want. The cost of the cloud factors in lots of stuff, not just hardware, factors such as management, infrastructure, cost of electricity, and so on. So it’s really hard to figure out what components you need to make the most efficient system.”

As part of his doctoral research, Conley had been focusing on new technologies before jumping on the cloud bandwagon. “We did do some experimentation with some very high-speed technologies that aren’t even available in the cloud today,” the Ph.D. student recalled. “We thought they would be present in maybe five years, because the research is forward thinking. But they cost a lot and we didn’t have the money to scale up. The main advantage of the cloud is that you don’t have to pay the capital expense upfront.”

Unfortunately, when the CNS team went looking to use relatively high-speed servers at Amazon EC2, the servers were in short supply. “They don’t have all that many of these servers available because we aren’t the standard customer of the cloud,” said CSE professor Porter. “So we ran into problems on the highway, even just getting enough lanes on the highway to drive.”

Switching metaphors, Porter added that “an all-you-can-eat buffet works pretty well for most people, but if the San Diego Chargers all pile out of a bus at the same time, you might run out of food.” That also means that users are limited in what they can get from cloud providers, and until now, clouds run by Amazon, Google and Microsoft have had plenty of business from larger cohorts of users without the narrow requirements of, for example, an academic center such as CNS. “The sorting work that we do stresses out their system in a way that they did not design it

for,” said Porter. “We wanted to make the argument that sorting, although it seems like a mundane topic, looks a lot like big-data processing, so if you want to do big-data processing, you have to be good at sorting. And the current cloud is not that good at sorting.”

Indeed, setting world records in data sorting requires a sustained amount of high performance over a short period of time, and all of the resources need to be working in concert for the entire duration of processing the big data. That’s difficult to do, because there are other people in the cloud trying to use the same resources at the same time. “The reason they added this CloudSort category to the competition this year was to draw people’s attention to the cloud, and to draw the attention of cloud providers to these records,” said Conley, who hopes to work in industry after earning his Ph.D. next March. “Having a CloudSort competition gives the cloud providers a metric to say, well, we should work on that and make it better.”

Academic researchers are also taking away lessons. “One of the fundamental surprises that came out of this research was that if you take the performance of a single server, it doesn’t extrapolate very well to figuring out the performance of a cluster of servers,” said Porter. “If you took the power of one server and multiplied it by 100, a 100-node cluster doesn’t perform as well. Maybe it performs 50 times as well, or 30 times.”

For the UC San Diego researchers, the big takeaway is efficiency, i.e., the fastest sorting across the fewest servers. The “EC” in Amazon EC2 stands for ‘elastic compute,’ and it was designed for users such as retailers during the holidays when they need to handle the load. “At Thanksgiving they can elastically extend use of the cloud for the holiday shopping season, then contract it after New Year’s,” noted Conley. “So the retailer doesn’t have to provision their system all year round.”

But increasingly, say Porter and Conley, large companies dealing with big data in the cloud are running into the same problems facing the UCSD researchers. Insisted Porter: “If you really talk to IT people and chief information officers, people like that worry about this.”

Meanwhile, cloud providers are gradually increasing the options available to customers (now counting 37 flavors at EC2), and those customers are shopping around. At UC San Diego, Porter and CNS acting director Stefan Savage are working with four students in the Emerging Research Scholars Program (ERSP) to undertake a comparison of Amazon EC2 versus its counterparts at Google and Microsoft.

Funding for the data-sorting project was provided by the National Science Foundation (NSF) through a grant to support highly-efficient, pipeline-oriented, data-intensive scalable computing (grant #1116079). The three-year grant ended in August.

MEDIA CONTACT

Doug Ramsey, , dramsey@ucsd.edu

UC San Diego's [Studio Ten 300](#) offers radio and television connections for media interviews with our faculty, which can be coordinated via studio@ucsd.edu. To connect with a UC San Diego faculty expert on relevant issues and trending news stories, visit <https://ucsdnews.ucsd.edu/media-resources/faculty-experts>.