


High-performance computing cuts particle collision data prep time

November 28 2017



This animation shows a series of collision events at STAR, each with thousands of particle tracks and the signals registered as some of those particles strike various detector components. It should give you an idea of how complex the challenge is to reconstruct a complete record of every single particle and the conditions under which it was created so scientists can compare hundreds of millions of events to look for trends and make discoveries. Credit: Brookhaven National Laboratory

For the first time, scientists have used high-performance computing (HPC) to reconstruct the data collected by a nuclear physics

experiment—an advance that could dramatically reduce the time it takes to make detailed data available for scientific discoveries.

The demonstration project used the Cori supercomputer at the National Energy Research Scientific Computing Center (NERSC), a high-performance computing center at Lawrence Berkeley National Laboratory in California, to reconstruct multiple datasets collected by the STAR detector during particle collisions at the Relativistic Heavy Ion Collider (RHIC), a nuclear physics research facility at Brookhaven National Laboratory in New York. By running multiple computing jobs simultaneously on the allotted supercomputing cores, the team transformed 4.73 petabytes of [raw data](#) into 2.45 petabytes of "physics-ready" data in a fraction of the time it would have taken using in-house high-throughput computing resources, even with a two-way transcontinental data journey.

"The reason why this is really fantastic," said Brookhaven physicist Jérôme Lauret, who manages STAR's computing needs, "is that these high-performance computing resources are elastic. You can call to reserve a large allotment of computing power when you need it—for example, just before a big conference when physicists are in a rush to present new results." According to Lauret, preparing raw data for analysis typically takes many months, making it nearly impossible to provide such short-term responsiveness. "But with HPC, perhaps you could condense that many months production time into a week. That would really empower the scientists!"

The accomplishment showcases the synergistic capabilities of RHIC and NERSC—U.S. Department of Energy (DOE) Office of Science User Facilities located at DOE-run national laboratories on opposite coasts—connected by one of the most extensive high-performance data-sharing networks in the world, DOE's Energy Sciences Network (ESnet), another DOE Office of Science User Facility.

"This is a key usage model of high-performance computing for experimental data, demonstrating that researchers can get their raw data processing or simulation campaigns done in a few days or weeks at a critical time instead of spreading out over months on their own dedicated resources," said Jeff Porter, a member of the data and analytics services team at NERSC.

Billions of data points

To make physics discoveries at RHIC, scientists must sort through hundreds of millions of collisions between ions accelerated to very high energy. STAR, a sophisticated, house-sized electronic instrument, records the subatomic debris streaming from these particle smashups. In the most energetic events, many thousands of particles strike detector components, producing firework-like displays of colorful particle tracks. But to figure out what these complex signals mean, and what they can tell us about the intriguing form of matter created in RHIC's collisions, scientists need detailed descriptions of all the particles and the conditions under which they were produced. They must also compare huge statistical samples from many different types of collision events.

Cataloging that information requires sophisticated algorithms and pattern recognition software to combine signals from the various readout electronics, and a seamless way to match that data with records of collision conditions. All the information must then be packaged in a way that physicists can use for their analyses.



Cori, the newest supercomputer at the National Energy Research Scientific Computing Center (NERSC), is a Cray XC40 with a peak performance of about 30 petaflops. Credit: Brookhaven National Laboratory

Since RHIC started running in the year 2000, this raw data processing, or reconstruction, has been carried out on dedicated computing resources at the RHIC and ATLAS Computing Facility (RACF) at Brookhaven. High-throughput computing (HTC) clusters crunch the data, event-by-event, and write out the coded details of each collision to a centralized mass storage space accessible to STAR physicists around the world.

But the challenge of keeping up with the data has grown with RHIC's ever-improving collision rates and as new detector components have been added. In recent years, STAR's annual raw data sets have reached

billions of events with data sizes in the multi-Petabyte range. So the STAR computing team investigated the use of external resources to meet the demand for timely access to physics-ready data.

Many cores make light work

Unlike the high-throughput computers at the RACF, which analyze events one-by-one, HPC resources like those at NERSC break large problems into smaller tasks that can run in parallel. So the first challenge was to "parallelize" the processing of STAR event data.

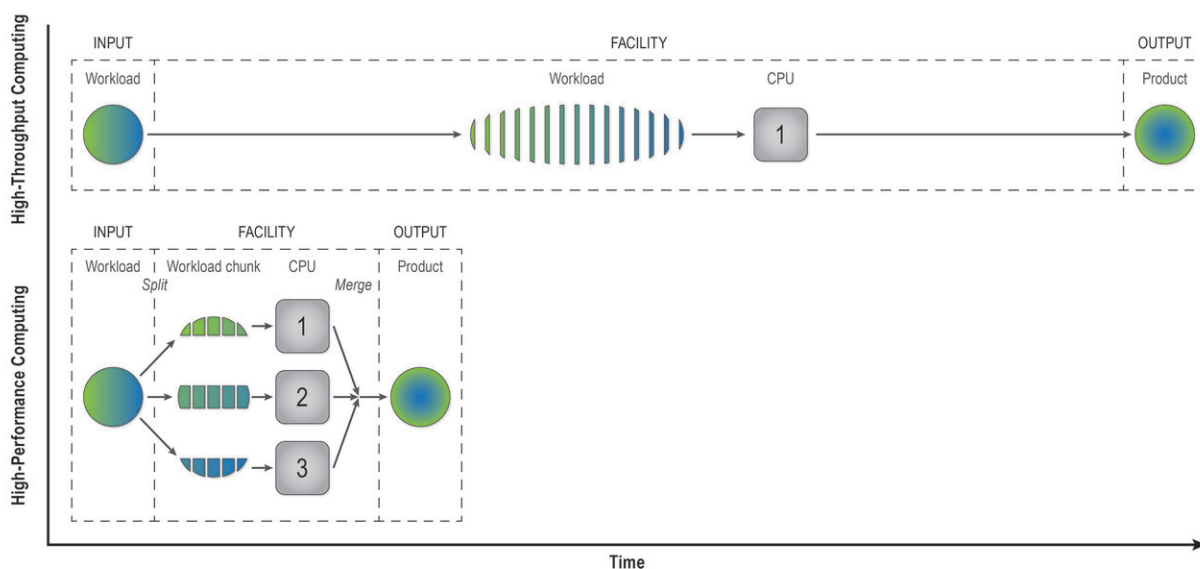
"We wrote workflow programs that achieved the first level of parallelization—event parallelization," Lauret said. That means they submit fewer jobs made of many events that can be processed simultaneously on the many HPC computing cores.

"Imagine building a city with 100 homes. If this was done in high-throughput fashion, each home would have one builder doing all the tasks in sequence—building the foundation, the walls, and so on," Lauret said. "But with HPC we change the paradigm. Instead of one worker per house we have 100 workers per house, and each worker has a task—building the walls or the roof. They work in parallel, at the same time, and we assemble everything together at the end. With this approach, we will build that house 100 times faster."

Of course, it takes some creativity to think about how such problems can be broken up into tasks that can run simultaneously instead of sequentially, Lauret added.

HPC also saves time matching raw detector signals with data on the environmental conditions during each event. To do this, the computers must access a "condition database"—a record of the voltage, temperature, pressure, and other detector conditions that must be

accounted for in understanding the behavior of the particles produced in each collision. In event-by-event, high-throughput reconstruction, the computers call up the database to retrieve data for every single event. But because HPC cores share some memory, events that occur close in time can use the same cached condition data. Fewer calls to the database means faster data processing.



In high-throughput computing, a workload made up of data from many STAR collisions is processed event-by-event in a sequential manner to give physicists "reconstructed data" —the product they need to fully analyze the data. High-performance computing breaks the workload into smaller chunks that can be run through separate CPUs to speed up the data reconstruction. In this simple illustration, breaking a workload of 15 events into three chunks of five events processed in parallel yields the same product in one-third the time as the high-throughput method. Using 32 CPUs on a supercomputer like Cori can greatly reduce the time it takes to transform the raw data from a real STAR dataset, with many millions of events, into useful information physicists can analyze to make discoveries. Credit: Brookhaven National Laboratory

Networking teamwork

Another challenge in migrating the task of raw data reconstruction to an HPC environment was just getting the data from New York to the supercomputers in California and back. Both the input and output datasets are huge. The team started small with a proof-of-principle experiment—just a few hundred jobs—to see how their new workflow programs would perform.

"We had a lot of assistance from the networking professionals at Brookhaven," said Lauret, "particularly Mark Lukasczyk, one of our network engineers, who was so excited about the science and helping us make discoveries." Colleagues in the RACF and ESnet also helped identify hardware issues and developed solutions as the team worked closely with Jeff Porter, Mustafa Mustafa, and others at NERSC to optimize the data transfer and the end-to-end workflow.

Start small, scale up

After fine-tuning their methods based on the initial tests, the team started scaling up to using 6,400 computing cores at NERSC, then up and up and up.

"6,400 cores is already half of the size of the resources available for data reconstruction at RACF," Lauret said. "Eventually we went to 25,600 cores in our most recent test." With everything ready ahead of time for an advance-reservation allotment of time on the Cori supercomputer, "we did this test for a few days and got an entire data production done in no time," Lauret said. According to Porter at NERSC, "This model is potentially quite transformative, and NERSC has worked to support such resource utilization by, for example, linking its center-wide high-performant disk system directly to its data transfer infrastructure and allowing significant flexibility in how job slots can be scheduled."

The end-to-end efficiency of the entire process—the time the program was running (not sitting idle, waiting for computing resources) multiplied by the efficiency of using the allotted supercomputing slots and getting useful output all the way back to Brookhaven—was 98 percent.

"We've proven that we can use the HPC resources efficiently to eliminate backlogs of unprocessed data and resolve temporary resource demands to speed up science discoveries," Lauret said.

He's now exploring ways to generalize the workflow to the Open Science Grid—a global consortium that aggregates [computing resources](#)—so the entire community of high-energy and nuclear physicists can make use of it.

More information: Mustafa Mustafa et al. STAR Data Reconstruction at NERSC/Cori, an adaptable Docker container approach for HPC, *Journal of Physics: Conference Series* (2017). [DOI: 10.1088/1742-6596/898/8/082023](#)

Provided by Brookhaven National Laboratory

Citation: High-performance computing cuts particle collision data prep time (2017, November 28) retrieved 23 June 2024 from <https://phys.org/news/2017-11-high-performance-particle-collision-prep.html>

This document is subject to copyright. Apart from any fair dealing for the purpose of private study or research, no part may be reproduced without the written permission. The content is provided for information purposes only.