

Berkeley Lab Technology Dramatically Speeds Up Searches of Large Databases

May 17 2005



In the world of physics, one of the most elusive events is the creation and detection of "quark-gluon plasma," the theorized atomic outcome of the "Big Bang" which could provide insight into the origins of the universe. By using experiments that involve millions of particle collisions, researchers hope to find unambiguous evidence of quarkgluon plasma.

Image: In this full-energy collision between gold ions at Brookhaven Lab's Relativistic Heavy Ion Collider, as captured by the STAR detector, the tracks indicate the paths taken by thousands of subatomic particles produced in the collisions as they pass through STAR's Time Projection



Chamber, a large, 3-D digitial camera, designed at Berkeley Lab.

Scientists describe such a collision with unambiguous evidence as a "rare event," which may be an understatement. For example, out of hundreds of millions of particle collisions in one experiment, an analysis found that only 80 collisions or "events" merited further study as scientists search for evidence of "jet quenching," a phenomenon that may indicate the existence of quark-gluon plasma. Other research into such exotic physics phenomena as "strangelets" needs to go through similar search processes.

Compounding the complexity of the search is the fact that the data files are on mass storage systems around the world, so locating and extracting these scientific needles from a virtual haystack of information would be very time-consuming and labor-intensive. For example, the brute-force approach of reading every record of the petabytes of distributed data from the Relativistic Heavy Ion Collider experiment called STAR at Brookhaven National Laboratory could take weeks at a time. The key to speeding up the searching process is to quickly locate those interesting events while ignoring millions of others so the important data can be extracted for further analysis.

Now, a search technology developed by researchers at the U.S. Department of Energy's Lawrence Berkeley National Laboratory makes the job much easier. The technology, known as the Word-Aligned Hybrid (WAH) compression method, was developed and recently patented by John Wu, Arie Shoshani and Ekow Otoo of Berkeley Lab's Scientific Data Management (SDM) Research Group.

The technique and its application are described in a paper recently selected as a "best paper" by the International Supercomputer Conference, and Wu will present the paper at the conference to be held June 21-24 in Heidelberg, Germany.



WAH is currently used in a software package called FastBit to compress bitmap indexes. A bitmap index is a method of reducing the response time of queries involving common types of conditions in data objects, such as "state = CA" and "age \geq 21." It achieves this by storing certain pre-computed answers as bitmaps. For example, a bitmap index for "state" might have one bitmap for each state in the U.S. Because computers can manipulate bitmaps efficiently, bitmap indices are efficient in searching for interesting records in large datasets.

WAH compression makes the bitmap index optimal in terms of computational complexity. A small number of the most efficient indexing schemes have this optimality property. What makes the new technology unique is that WAH-compressed indexes significantly outperform other schemes in tests.

"In tests conducted using actual data from high-energy physics experiments, we confirmed that our FastBit software is an order of magnitude faster than the best-known bitmap indexing schemes on average," according to Wu, the lead developer of FastBit.

Thanks to work led by Berkeley Lab, the physicists working on the STAR (Solenoidal Tracker at RHIC) high-energy physics experiment at Brookhaven now have a much more efficient tool in their search for evidence of the quark-gluon plasma. As their research evolves and the complexity of the problem unravels, scientists are finding that a new state of matter was definitely created at STAR, but to unambiguously characterize this new state of matter as the quark-gluon plasma, they need more sophisticated search criteria to locate the "rare" collision events that would contain the clear signatures of the plasma.

Grid Collector, the software module for the STAR analysis framework, uses two technologies to provide STAR analysts with a new way of accessing collision data. The first is FastBit's searching capability, and



the second is Storage Resource Managers (SRMs), which provide access to files stored on remote storage systems. Both technologies were developed as part of DOE's Scientific Discovery through Advanced Computing (SciDAC) Program. Instead of selecting the data files that contain the desired events as was previously done, analysts can now select events based on physically meaningful attributes known as tags. Through Grid Collector, analysis programs only read the selected events, instead of every event in the selected data files. Since most analysis jobs use only a fraction of the events in data files, the Grid Collector can significantly improve the turnaround time.

Without Grid Collector, many analysis jobs involving searches for rare events were considered nearly unfeasible. For example, Markus Oldenburg of Berkeley Lab's Nuclear Science Division and his colleagues were interested in 80 special events collected in 2001. Most participants in the project thought they could make more progress by pursuing alternative signatures, rather than spending the time to extract these 80 events. With Grid Collector, the researchers were able to extract the events in 15 minutes.

"The Grid Collector has opened new avenues for many challenging analysis jobs that we had to ignore or delay. These jobs are now practical with this innovative technology," said Jerome Lauret, software coordinator for the RHIC/STAR experiment. "By using FastBit, we may have very well abolished one limiting factor for our field."

Creating fast, manageable indexes

Indexing methods are used extensively by database management systems to provide fast query processing for users. While an indexing method can make it easier to search the data, indexes themselves — especially bitmap indexes — can require a larger amount of additional storage space. If the index becomes too large, it's unusable. One way to address



this issue is to compress the indexes. However, this may also increase the time needed to perform search operations. A number of specialized compression schemes have been proposed to process compressed indexes efficiently, with the best-known one called the Byte-aligned Bitmap Code (BBC).

The goal of the Berkeley Lab project was to create an indexing system that could be compressed and at the same time offers much faster searches than existing methods. To achieve this goal, the WAH compression scheme was developed. While WAH-compressed indexes are slightly larger than BBC-compressed indexes, the time needed to process a query is less, often much less.

"We were seeking a worthwhile space-time tradeoff and we succeeded," Wu said. "What makes our compressed bitmap index really special is that it is not only theoretically optimal but also practically more efficient than any other indexing scheme tested."

This new technology, officially called "Word Aligned Bitmap Compression Method and Data Structure," is currently being used by other DOE research projects and has yielded several success stories. Tracking features in the analysis of combustion simulation data is more efficient. By using FastBit and compressed bitmaps, the FastBit team was able to significantly speed up the problem of tracking ignition kernels in a hydrogen-air combustion simulation. This approach addressed the difficult problem of identifying multi-attribute data distinguishing the ignition kernels from the rest of the simulation data and tracking the progression of flames over time. This was done in collaboration with Wendy Koegler and Jacqueline Chen of Sandia National Laboratories.

DEX, or Dexterous Data Explorer, is a collaboration between the Scientific Data Management Group and Berkeley Lab's Visualization



Group. DEX uses FastBit to provide query-based visualization of large scientific datasets. A preliminary version of the software was demonstrated at the Supercomputing 2004 conference on both combustion datasets and supernova simulation datasets. Berkeley Lab collaborators on DEX are Kurt Stockinger, John Shalf and Wes Bethel.

Compressed bitmaps are also used in view-dependent isosurface software. At Supercomputing 2004, a preliminary version of the software was demonstrated to display in real time the isosurfaces of large complex data produced from a simulation of the Rayleigh-Taylor instability in computational fluid dynamics. In this application, compressed bitmaps are used to record what data items are visible from a particular viewing angle. This allows the software to extract the minimal amount of data items required for visualization and to render in real time very large complex isosurfaces as the user changes viewing angles. This was done in collaboration with Guruprasad Kora, Jian Huang, Jinzhu Gao and Nagiza Samatova of Oak Ridge National Laboratory.

The effectiveness of FastBit has attracted the attention of other institutions as well. At CERN, the developers of ROOT, an objectoriented data analysis framework, have started evaluating the incorporation of FastBit into their software. Since the ROOT software is used by most major high-energy physics projects around the world, fully integrating FastBit into ROOT would make the efficient search capability of FastBit available to a large user community.

"We have even learned that a Brazilian telecommunications provider has implemented WAH compression for their data analysis," Wu said.

Source: Berkeley Lab



Citation: Berkeley Lab Technology Dramatically Speeds Up Searches of Large Databases (2005, May 17) retrieved 30 April 2024 from <u>https://phys.org/news/2005-05-berkeley-lab-technology-large-databases.html</u>

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.