

Petacache: Use that Memory

March 7 2006



SLAC's Computerraum

For decades, high energy experimental physicists have struggled with a fundamental problem: they simply have too much data to analyze quickly and in its entirety.

BaBar researchers routinely wait nine months for computers to sift through large datasets, searching for interesting events and setting these aside for later analysis. This “data skimming” alone constantly uses about 50 percent of BaBar's computing power. And that’s before a researcher can even start analyzing her or his data. Preparing data from CERN's Large Hadron Collider (LHC) will only take longer.

Recognizing this widespread limitation, a team at SLAC is developing the PetaCache project, a new way of thinking about data access and storage. With new computer software and more efficient types of

memory, PetaCache may significantly increase the speed of data analysis.

"PetaCache may help scientists change the way they think about exploring new ideas," said PetaCache project manager Randal Melen. "It will allow a physicist with a sudden new idea, an 'I wonder if...' moment, to quickly begin exploring that new idea."

Before the early 1990s, researchers analyzed much of their data from magnetic tape, having their computers spool through miles of it to find interesting events. As disk drives got larger and cheaper, and with the rise of computer clusters, much more of the data could be kept on disk. Yet these disks still required mechanical movement, limiting the speed at which researchers could begin accessing data. Computer technology has made great strides in speeding up the movement of data—called bandwidth—but the time to get the first byte of data—called latency—has been much slower to improve. "PetaCache, then, is really about improving the latency of testing new ideas," said Melen.

To do this, PetaCache uses several types of memory, not disks. Although memory is much faster at getting this first byte of data, in the past it has been too expensive to buy in the quantities necessary to record and analyze the massive amounts of data taken at particle accelerators. Today, DRAM (Dynamic Random Access Memory) and flash memory are more affordable, and flash memory is expected to continue to drop in price as it is used more and more in consumer electronics such as digital cameras, iPod-like devices, and cell phones. If successful, the PetaCache project will allow researchers to use both DRAM and flash memory on a large scale.

The prototype PetaCache system comprises two racks of 64 server computers, each with 16 gigabytes of DRAM for a total of one terabyte of memory. This large yet fragmented amount of memory is linked

together with SCALLA (Structured Cluster Architecture for Low Latency Access), a computer program developed by SCCS Software Developer Andy Hanushevsky. SCALLA moves data from data servers to batch systems running physics analysis software with the lowest possible latencies. This load-balancing, self-organizing software distributes data across many data servers efficiently, making the individual machines appear as one huge chunk of memory to SCALLA-aware physics applications.

"The software makes good use of common hardware, so you don't have to make huge expenditures for great computing power," said Hanushevsky.

Right now, SLAC's prototype system has one terabyte (1,000 gigabytes) of DRAM memory. With their next machine, the PetaCache team hopes to mainly use less expensive flash memory which, according to SCCS Director Richard Mount, "holds future promise of cost-effective memory-based data-analysis systems."

This second-generation prototype will aim at a few tens of terabytes of flash memory, which would make the system useful to BaBar and LSST researchers. In the next decade, the PetaCache team hopes to expand the system to a petabyte (1,000 terabytes). This is around the scale of what is needed to be useful at the LHC.

"Over the next few years, this type of memory technology will become much more common, from BaBar to the LHC to banks and airline reservation systems," said Research Director Emeritus David Leith. "They all benefit from being able to work from memory."

Source: Stanford Linear Accelerator Center, by Kelen Tuttle

Citation: Petacache: Use that Memory (2006, March 7) retrieved 25 April 2024 from <https://phys.org/news/2006-03-petacache-memory.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.