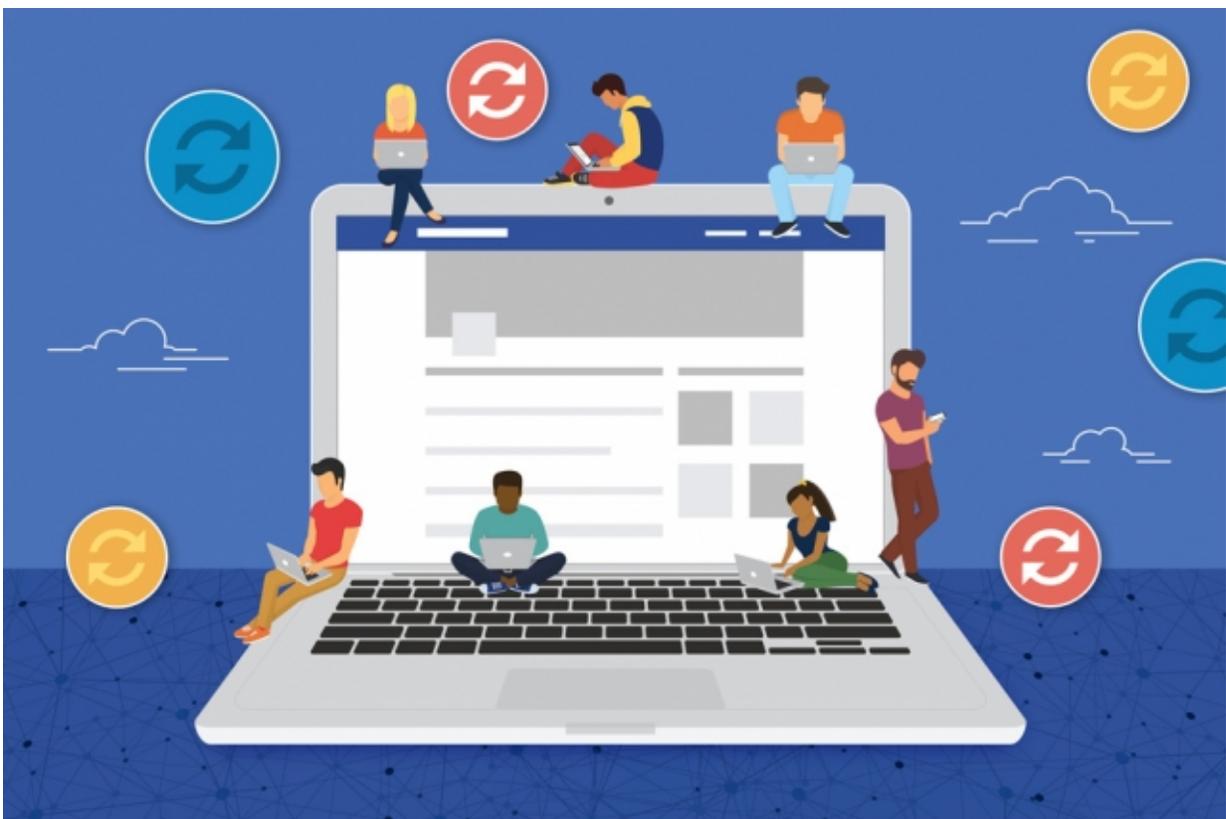


Flash-memory system could reduce power consumption of data center 'caches' by 90 percent

August 31 2017, by Larry Hardesty



Researchers from CSAIL have devised a new system for data center caching that uses flash memory. In addition to costing less and consuming less power, a flash caching system could dramatically reduce the number of cache servers required by a data center. Credit: MIT News

Most modern websites store data in databases, and since database queries are relatively slow, most sites also maintain so-called cache servers, which list the results of common queries for faster access. A data center for a major web service such as Google or Facebook might have as many as 1,000 servers dedicated just to caching.

Cache servers generally use random-access memory (RAM), which is fast but expensive and power-hungry. This week, at the International Conference on Very Large Databases, researchers from MIT's Computer Science and Artificial Intelligence Laboratory (CSAIL) are presenting a new system for [data center](#) caching that instead uses [flash memory](#), the kind of memory used in most smartphones.

Per gigabyte of memory, [flash](#) consumes about 5 percent as much energy as RAM and costs about one-tenth as much. It also has about 100 times the storage density, meaning that more data can be crammed into a smaller space. In addition to costing less and consuming less power, a flash caching system could dramatically reduce the number of cache servers required by a data center.

The drawback to flash is that it's much slower than RAM. "That's where the disbelief comes in," says Arvind, the Charles and Jennifer Johnson Professor in Computer Science Engineering and senior author on the conference paper. "People say, 'Really? You can do this with flash memory?' Access time in flash is 10,000 times longer than in DRAM [dynamic RAM]."

But slow as it is relative to DRAM, flash access is still much faster than human reactions to new sensory stimuli. Users won't notice the difference between a request that takes .0002 seconds to process—a typical round-trip travel time over the internet—and one that takes .0004 seconds because it involves a flash query.

Keeping pace

The more important concern is keeping up with the requests flooding the data center. The CSAIL researchers' system, dubbed BlueCache, does that by using the common computer science technique of "pipelining." Before a flash-based cache server returns the result of the first query to reach it, it can begin executing the next 10,000 queries. The first query might take 200 microseconds to process, but the responses to the succeeding ones will emerge at .02-microsecond intervals.

Even using pipelining, however, the CSAIL researchers had to deploy some clever engineering tricks to make flash caching competitive with DRAM caching. In tests, they compared BlueCache to what might be called the default implementation of a flash-based cache server, which is simply a data-center database server configured for caching. (Although slow compared to DRAM, flash is much faster than magnetic hard drives, which it has all but replaced in data centers.) BlueCache was 4.2 times as fast as the default implementation.

Joining Arvind on the paper are first author Shuotao Xu and his fellow MIT graduate student in electrical engineering and computer science Sang-Woo Jun; Ming Liu, who was an MIT graduate student when the work was done and is now at Microsoft Research; Sungjin Lee, an assistant professor of computer science and engineering at the Daegu Gyeongbuk Institute of Science and Technology in Korea, who worked on the project as a postdoc in Arvind's lab; and Jamey Hicks, a freelance software architect and MIT affiliate who runs the software consultancy Accelerated Tech.

The researchers' first trick is to add a little DRAM to every BlueCache flash cache—a few megabytes per million megabytes of flash. The DRAM stores a table which pairs a database query with the flash-memory address of the corresponding query result. That doesn't make

cache lookups any faster, but it makes the detection of cache misses—the identification of data not yet imported into the cache—much more efficient.

That little bit of DRAM doesn't compromise the system's energy savings. Indeed, because of all of its added efficiencies, BlueCache consumes only 4 percent as much power as the default implementation.

Engineered efficiencies

Ordinarily, a cache system has only three operations: reading a value from the cache, writing a new value to the cache, and deleting a value from the [cache](#). Rather than rely on software to execute these operations, as the default implementation does, Xu developed a special-purpose hardware circuit for each of them, increasing speed and lowering power consumption.

Inside a BlueCache server, the flash memory is connected to the central processor by a wire known as a "bus," which, like any data connection, has a maximum capacity. BlueCache amasses enough queries to exhaust that capacity before sending them to memory, ensuring that the system is always using communication bandwidth as efficiently as possible.

With all these optimizations, BlueCache is able to perform write operations as efficiently as a DRAM-based system. Provided that each of the query results it's retrieving is at least eight kilobytes, it's as efficient at read operations, as well. (Because flash memory returns at least eight kilobytes of data for any request, it's efficiency falls off for really small query results.)

BlueCache, like most data-center caching systems, is a so-called key-value store, or KV store. In this case, the key is the database query and the value is the response.

"The flash-based KV store architecture developed by Arvind and his MIT team resolves many of the issues that limit the ability of today's enterprise systems to harness the full potential of flash," says Vijay Balakrishnan, director of the Data Center Performance and Ecosystem program at Samsung Semiconductor's Memory Solutions Lab. "The viability of this type of system extends beyond caching, since many data-intensive applications use a KV-based software stack, which the MIT team has proven can now be eliminated. By integrating programmable chips with flash and rewriting the software stack, they have demonstrated that a fully scalable, performance-enhancing storage technology, like the one described in the paper, can greatly improve upon prevailing architectures."

This story is republished courtesy of MIT News (web.mit.edu/newsoffice/), a popular site that covers news about MIT research, innovation and teaching.

Provided by Massachusetts Institute of Technology

Citation: Flash-memory system could reduce power consumption of data center 'caches' by 90 percent (2017, August 31) retrieved 19 April 2024 from <https://techxplore.com/news/2017-08-flash-memory-power-consumption-center-caches.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.