

Innovation promises to cut massive power use at big data companies in a flash

July 19 2012



Vivek Pai (above), an associate professor of computer science at Princeton, worked with graduate student Anirudh Badam to develop a software technique that could radically cut power use. Credit: Photo by Frank Wojciechowski

Big data needs big power. The server farms that undergird the Internet run on a vast tide of electricity. Even companies that have invested in upgrades to minimize their eco-footprint use tremendous amounts: The New York Times estimates that Google, for example, uses enough electricity in its data centers to power about 200,000 homes.

Now, a team of Princeton University engineers has a solution that could radically cut that power use. Through a new software technique, researchers from the School of Engineering and [Applied Science](#) have opened the door for companies to use a new type of memory in their

servers that demands far less energy than the current systems.

The software, called SSDAlloc, allows the companies to substitute solid state memory, commonly called flash memory, for the more expensive and energy-intensive type of memory that is now used for most computer operations.

"The biggest potential users are the big data centers," said Vivek Pai, an associate professor of computer science who developed the program with his graduate student, Anirudh Badam. "They are going to see the greatest improvements."

A version of SSDAlloc is already being used with high-end flash memory manufactured by Fusion-io, of [Salt Lake City](#). Princeton has signed a non-exclusive [licensing agreement](#) with the company. Brent Compton, the Fusion-io's senior director of product management, said the software "simplifies performance for developers in ways that were out of reach just a couple of years ago."

The massive server centers that support operations ranging from online shopping to social media are built around a type of [computer memory](#) called [random access memory](#), or RAM. While very fast and flexible, RAM needs a constant stream of electricity to operate.

The power not only costs money, it also generates heat that forces the companies to spend more funds on cooling.

The Princeton engineers' program allows the data companies to substitute flash memory — similar to chips used in "thumb drives" — for much of their RAM. Unlike RAM, flash only uses small amounts of electricity, so switching memory types can drastically cut a company's power bill. In extreme cases, depending on the type of programs run by the servers, that reduction can be as much as 90 percent (compared to a

computer using RAM alone). And because those machines are not generating as much heat, the data centers can also cut their cooling bills.

Flash memory is also about 10 times cheaper than RAM, so companies can also save money on hardware upfront, Pai said.

So how does it work?

Badam, a graduate student in computer science, said that SSDAlloc basically changes the way that programs look for data in a computer.

Traditionally, a computer program will run its operations in RAM, which is fast and efficient, but unable to store information without power. When the program needs to store information longer, or when it needs to use data that is not in the RAM, it looks to storage memory — either flash memory or mechanical hard drives.

That is where a bottleneck occurs. The step at which the program switches to storage memory is glacially slow in computer terms. That is often the nature of the storage medium itself — mechanical hard drives are vastly slower than RAM. But it is also the result of underlying operating systems, such as Linux or Windows, that govern how the computer searches for information.

Flash memory is much faster than a hard drive, and flash is getting faster all the time. Currently, high-end flash memory has retrieval speeds of a million requests per second. (A top mechanical hard drive's retrieval speed is about 300 requests per second.)

That discrepancy created a dilemma for flash. The physical flash memory itself was fast enough that it could operate as an extension of RAM, but the underlying retrieval system's slow speed throttled its potential.

Based on earlier research, funded in part by the National Science Foundation, Pai and Badam felt they had a technique that could universally allow flash memory to serve as an extension of RAM. Other researchers had developed more narrow techniques, but they were difficult to program and only worked for certain applications. The Princeton researchers' idea would work with any program with minimal, and relatively straightforward, adjustments.

It was an ambitious idea. Two other research teams outside of Princeton had tried unsuccessfully to create similar results, and many experts were convinced that the technique could not be done through changes in software alone.

"It did seem like a long shot," Badam said.

What Badam did was write software that allows programmers to bypass this traditional system of searching for information in storage memory. His system allows for requests for information that take advantage of flash memory's extremely fast retrieval times. Essentially, SSDAlloc moves the flash memory up in the internal hierarchy of computer data — instead of thinking of flash as a version of a storage drive, SSDAlloc tells the computer to consider it a larger, somewhat slower, version of RAM.

"I wanted to make [flash memory](#) look like it was traditional memory," he said.

The first version of the software required programmers to write a very small percentage of their software — Pai estimates about 1 percent — to work with SSDAlloc. But while completing a scientific internship at Fusion-io last summer, Badam was able to refine SSDAlloc so that programmers no longer have to alter any of their code to work with the system.

"A good thing about SSDAlloc is that it does not alter the program," Badam said. "If you were using RAM and you want to use RAM, you can do that. If you want to use solid state you can use that."

Pai predicts that the need for faster memory access will continue to grow as more computing relies on the virtual cloud rather than on individual machines. The cloud, of course, must be supported by servers running those programs.

"Our system monitors what the host system is doing and moves it into and out of [RAM](#) automatically," he said. "There is a whole class of applications in which this would be used."

Provided by Princeton University

Citation: Innovation promises to cut massive power use at big data companies in a flash (2012, July 19) retrieved 26 April 2024 from <https://phys.org/news/2012-07-massive-power-big-companies.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.