


# Computer engineers boost app speeds by more than nine percent

26 September 2016



Researchers from North Carolina State University and Samsung Electronics have found a way to boost the speed of computer applications by more than 9 percent. The improvement results from techniques that allow computer processors to retrieve data more efficiently.

Computer processors have to retrieve data from memory to perform operations. All data is stored in off-chip "main" memory. But data that the [processor](#) will use a lot is also stored – temporarily – in a die-stacked dynamic random access memory (DRAM) cache that is located closer to the processor, where it can be retrieved more quickly.

The data in the cache is organized into large blocks, or macroblocks, so that the processor knows where to find whatever data it needs. However, for any given operation, the processor doesn't need all of the data in a macroblock – and retrieving the unnecessary data takes time and energy.

To make the process more efficient, researchers have developed a technique in which the cache learns over time which data the processor needs from each macroblock. This allows the cache to do two things. First, the cache can compress the macroblock, retrieving only the relevant data. This

enables the cache to send data to the processor more efficiently. Second, because the macroblock is compressed, this frees up space in the cache that can be used to store other data, which the processor is more likely to need.

The researchers tested this approach, called Dense Footprint Cache, in a processor and memory simulator. After running 3 billion instructions for each application tested through the simulator, the researchers found that the Dense Footprint Cache sped up applications by 9.5 percent compared to state-of-the-art competing methods for managing die-stacked DRAM. Dense Footprint Cache also used 4.3 percent less energy.

The researchers also found that Dense Footprint Cache led to a significant improvement in "last-level cache miss ratios." Last-level cache misses occur when the processor tries to retrieve data from the cache, but the data aren't there, forcing the processor to retrieve the [data](#) from off-chip main memory. These cache misses make operations much less efficient – and Dense Footprint Cache reduced last-level cache miss ratios by 43 percent.

The work is featured in a paper, "[Dense Footprint Cache: Capacity-Efficient Die-Stacked DRAM Last Level Cache](#)," that will be presented at the International Symposium on Memory Systems, Oct. 3-6 in Washington, D.C.

Lead author of the paper is Seunghee Shin, a Ph.D. student at NC State. The paper was co-authored by Yan Solihin, a professor of electrical and computer engineering at NC State, and Sihong Kim of Samsung Electronics.

Provided by North Carolina State University

APA citation: Computer engineers boost app speeds by more than nine percent (2016, September 26)  
retrieved 25 June 2019 from <https://phys.org/news/2016-09-boost-app-percent.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.*