

New 'Performance Cloning' Techniques Designed to Boost Computer Chip Memory Systems Design

For Immediate Release

September 30, 2015

Yan Solihin (<mailto:solihin@ncsu.edu>) | 919.513.2965 (tel:919.513.2965)

Matt Shipman (mailto:matt_shipman@ncsu.edu) | 919.515.6386 (tel:919.515.6386)

North Carolina State University researchers have developed software using two new techniques to help computer chip designers improve memory systems. The techniques rely on "performance cloning," which can assess the behavior of software without compromising privileged data or proprietary computer code.

Computer chip manufacturers try to design their chips to provide the best possible performance. But to find the most effective designs, manufacturers need to know what sort of software their clients will be using.

"For example, programs that model protein folding use a lot of computing power, but very little data – so manufacturers know to design chips with lots of central processing units (CPUs), but significantly less memory storage than would be found on other chips," says Yan Solihin, an associate professor of electrical and computer engineering at NC State and an author of two papers describing the new techniques.

However, many large customers – from major corporations to Wall Street firms – don't want to share their code with outsiders. And that makes it tough for chip manufacturers to develop the best possible chip designs.

One way to address this problem is through performance cloning. The concept behind performance cloning is that a chip manufacturer would give profiler software to a client. The client would use the profiler to assess its proprietary software, and the profiler would then generate a statistical report on the proprietary software's performance. That report could be given to the chip manufacturer without compromising the client's data or code.

The profiler report would then be fed into generator software, which can develop a synthetic program that mimics the performance characteristics of the client's software. This synthetic program would then serve as the basis for designing chips that will better meet the client's needs.

Previous work at Ghent University and the University of Texas at Austin has used performance cloning to address issues related to CPU design – but those initiatives did not focus on memory systems, which are an important element of overall chip design.

Researchers have now developed software using two new techniques to help optimize memory systems.

The first technique, called MEMST (Memory EMulation using Stochastic Traces), assesses memory in a synthetic program by focusing on the amount of memory a program uses, the location of the data being retrieved and the pattern of retrieval.

For example, MEMST looks at how often a program retrieves data from the same location in a short period of time, and at how likely a program is to retrieve data from a location that is near other data that's been retrieved recently. Both of these variables affect how quickly the program can retrieve data.

The second technique, called MeToo, focuses on memory timing behavior – how often the program retrieves data and whether the program has periods in which it makes many memory requests in a short time. Memory timing behavior can have a significant impact on how a system's memory system is designed.

For example, if you think of memory requests as cars, you don't want to have a traffic jam – so you may want to be sure there are enough lanes for the traffic. These traffic lanes equate to memory bandwidth; the broader the bandwidth, the more lanes there are.

"Both MEMST and MeToo are useful for chip designers, particularly for designers who work on memory components, such as DRAM, memory controllers and memory buses," Solihin says.

 ([http://www.facebook.com/sharer/sharer.php?](http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fnews.ncsu.edu%2F2015%2F09%2Fsolihin-text=)

The new techniques expand on previous work done by Solihin that used performance cloning to look at cache memory.

"Our next step is to take MEMST and MeToo, as well as our work on cache memory, and develop an integrated program that we can commercialize," says Solihin, author of the forthcoming [Fundamentals of Parallel Multicore Architecture](https://www.crcpress.com/Fundamentals-of-Parallel-Multicore-Architecture/Solihin/9781482211184) (<https://www.crcpress.com/Fundamentals-of-Parallel-Multicore-Architecture/Solihin/9781482211184>), which addresses memory hierarchy design.

The paper on MEMST, "[MEMST: Cloning Memory Behavior using Stochastic Traces](https://drive.google.com/file/d/0B7MY_IMKt1T-X0ZKbTBfOWQ5M2c/view) (https://drive.google.com/file/d/0B7MY_IMKt1T-X0ZKbTBfOWQ5M2c/view)," will be presented at the International Symposium on Memory Systems, being held Oct. 5-8 in Washington, D.C. The paper was co-authored by Solihin and Ganesh Balakrishnan of Advanced Micro Devices, a former NC State Ph.D. student.

The paper on MeToo, "[MeToo: Stochastic Modeling of Memory Traffic Timing Behavior](https://drive.google.com/file/d/0B7MY_IMKt1T-ZHBzdVhDN1JPOFE/view) (https://drive.google.com/file/d/0B7MY_IMKt1T-ZHBzdVhDN1JPOFE/view)," will be presented at the International Conference on Parallel Architecture and Compilation, being held Oct. 18-21 in San Francisco, Calif. Lead author of the paper is Yipeng Wang, a Ph.D. student at NC State. Co-authors are Balakrishnan and Solihin. The work was supported by the National Science Foundation under grant number CNS-0834664.

-shipman-

Note to Editors: The study abstracts follow.

"MEMST: Cloning Memory Behavior Using Stochastic Traces"

Authors: Ganesh Balakrishnan, Advanced Micro Devices; Yan Solihin, North Carolina State University

Presented: International Symposium on Memory Systems, Oct. 5-8, Washington, D.C.

Abstract: Memory Controller and DRAM architecture are critical aspects of Chip Multi Processor (CMP) design. A good design needs an in-depth understanding of end-user workloads. However, designers rarely get insights into end-user workloads because of the proprietary nature of source code or data. Workload cloning is an emerging approach that can bridge this gap by creating a proxy for the proprietary workload (clone). Cloning involves profiling workloads to glean key statistics and then generating a clone offline for use in the design environment. However, there are no existing cloning techniques for accurately capturing memory controller and DRAM behavior that can be used by designers for a wide design space exploration. We propose Memory EMulation using Stochastic Traces, MEMST, a highly accurate black box cloning framework for capturing DRAM and MC behavior. We provide a detailed analysis of statistics that are necessary to model a workload accurately. We will also show how a clone can be generated from these statistics using a novel stochastic method. Finally, we will validate our framework across a wide design space by varying DRAM organization, address mapping, DRAM frequency, page policy, scheduling policy, input bus bandwidth, chipset latency, DRAM die revision, DRAM generation and DRAM refresh policy. We evaluated MEMST using CPU2006, BioBench, Stream and PARSEC benchmark suites across the design space for single-core, dual-core, quad-core and octa-core CMPs. We measured both performance and power metrics for the original workload and clones. The clones show a very high degree of correlation with the original workload for over 7900 data points with an average error of 1.8% and 1.6% for transaction latency and DRAM power respectively.

"MeToo: Stochastic Modeling of Memory Traffic Timing Behavior"

Authors: Yipeng Wang and Yan Solihin, North Carolina State University; Ganesh Balakrishnan, Advanced Micro Devices

Presented: International Conference on Parallel Architecture and Compilation, Oct. 18-21, San Francisco, California

Abstract: The memory subsystem (memory controller, bus, and DRAM) is becoming a bottleneck in computer system performance. Optimizing the design of the multicore memory subsystem requires good understanding of the representative workload. A common practice in designing the memory subsystem is to rely on trace simulation. However, the conventional method of relying on traditional traces faces two major challenges. First, many software users are apprehensive about sharing their code (source or binaries) due to the proprietary nature of the code or secrecy of data, so representative traces are sometimes not available. Second, there is a feedback loop where memory performance affects processor performance, which in turn alters the timing of memory requests that reach the bus. Such feedback loop is difficult to capture with traces. In this paper, we present MeToo, a framework for generating synthetic memory traffic for memory subsystem design exploration. MeToo uses a small set of statistics that summarizes the performance behavior of the original applications, and generates synthetic traces or executables stochastically, allowing applications to remain proprietary. MeToo uses novel methods for mimicking the memory feedback loop. We validate MeToo clones, and show very good fit with the original applications' behavior, with an average error of only 4.2%, which is a small fraction of the errors obtained using geometric inter-arrival (commonly used in queueing models) and uniform inter-arrival.

SHARE

(<http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fnews.ncsu.edu%2F2015%2F09%2Fsolihin-memory-2015-09/>)

f (<http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fnews.ncsu.edu%2F2015%2F09%2Fsolihin-memory-2015%2F>)

t (<https://twitter.com/intent/tweet?text=New+%E2%80%99Performance+Cloning%E2%80%99+Techniques+Designed+to+Boost+>)

in ([http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fnews.ncsu.edu%2F2015%2F09%2Fsolihin-memory-2015%2F&title=New 'Performance Cloning' Techniques Designed to Boost Computer Chip Memory Systems Design&source='NC State News'\)](http://www.linkedin.com/shareArticle?mini=true&url=https%3A%2F%2Fnews.ncsu.edu%2F2015%2F09%2Fsolihin-memory-2015%2F&title=New+%E2%80%99Performance+Cloning%E2%80%99+Techniques+Designed+to+Boost+Computer+Chip+Memory+Systems+Design&source='NC+State+News'))

g+ (<https://plus.google.com/share?url=https%3A%2F%2Fnews.ncsu.edu%2F2015%2F09%2Fsolihin-memory-2015%2F>)

✉ ([mailto:?subject=New 'Performance Cloning' Techniques Designed to Boost Computer Chip Memory Systems Design&body=Computer science researchers have developed software using two new techniques to help computer chip designers improve memory systems. The techniques rely on "performance cloning," which can assess the behavior of software without compromising privileged data or proprietary computer code.

https://news.ncsu.edu/2015/09/solihin-memory-2015/](mailto:?subject=New+%E2%80%99Performance+Cloning%E2%80%99+Techniques+Designed+to+Boost+Computer+Chip+Memory+Systems+Design&body=Computer+science+researchers+have+developed+software+using+two+new+techniques+to+help+computer+chip+designers+improve+memory+systems.+The+techniques+rely+on+%E2%80%9Cperformance+cloning,%E2%80%9D+which+can+assess+the+behavior+of+software+without+compromising+privileged+data+or+proprietary+computer+code.

https://news.ncsu.edu/2015/09/solihin-memory-2015/))

FILED UNDER

Research and Innovation (<https://news.ncsu.edu/category/research-and-innovation/>)

TAGS

electrical and computer engineering (<https://news.ncsu.edu/tag/electrical-and-computer-engineering/>), engineering (<https://news.ncsu.edu/tag/engineering/>), news releases (<https://news.ncsu.edu/tag/news-releases/>), research news (<https://news.ncsu.edu/tag/research-news/>)

3 RESPONSES ON "NEW 'PERFORMANCE CLONING' TECHNIQUES DESIGNED TO BOOST COMPUTER CHIP MEMORY SYSTEMS DESIGN"

Pingback: Semiconductor Engineering .. System Bits: Oct. 6 (<http://semiengineering.com/system-bits-oct-6/>)

Caleb Kumar says: October 28, 2015 at 7:50 pm (<https://news.ncsu.edu/2015/09/solihin-memory-2015/#comment-7364501>)

Hi,

Very interesting article.....

I am actual going to present a summary of the posted two papers on the technique. I was wondering if I could get your opinion on the potential impact of these new techniques and/or just any other extra thoughts you had on the topic.

Thanks!

[Reply \(https://news.ncsu.edu/2015/09/solihin-memory-2015/?replytocom=7364501#respond\)](https://news.ncsu.edu/2015/09/solihin-memory-2015/?replytocom=7364501#respond)

Caleb Kumar says: October 28, 2015 at 8:23 pm (<https://news.ncsu.edu/2015/09/solihin-memory-2015/#comment-7364571>)

One more question.....

Are there any pros and cons to each technique? Is one more applicable to a piece of software than another?

Thanks!

[Reply \(https://news.ncsu.edu/2015/09/solihin-memory-2015/?replytocom=7364571#respond\)](https://news.ncsu.edu/2015/09/solihin-memory-2015/?replytocom=7364571#respond)

LEAVE A RESPONSE

Your email address will not be published. All fields are required.

Your Response

Name

Email

Submit Response

(<http://www.facebook.com/sharer/sharer.php?u=https%3A%2F%2Fnews.ncsu.edu%2F2015%2F09%2Fsolihin-memory-2015%2F>)

OTHER TOP NEWS

Hartman's Legacy of Pride

Damon Hartman's 56-yard field goal that secured NC State's win over Carolina in 1990 is the stuff of Wolfpack legend. His daughter, Hayden, a student member of the football team's sports medicine staff, gets a sideline view of the rivalry at Saturday's game at Kenan Stadium. →
(<https://news.ncsu.edu/2018/11/legacy-of-pride/>)

NC State Joins National Effort to Boost College Access

NC State joins an effort spearheaded by the Association of Public and Land-grant Universities to improve college access and student success. →
(<https://news.ncsu.edu/2018/11/powered-by-publics/>)

VIEW THE ARCHIVE →

([HTTPS://NEWS.NCSU.EDU/2018](https://news.ncsu.edu/2018/))