GEORGE MASON UNIVERSITY

DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING

Directed Speculation in Multi-core Memory Systems to Improve Performance and Efficiency

Speaker: Professor Paul Gratz

Texas A&M University

Wednesday, March 23, 2016

11am- 12pm, VSE Room 1602

Abstract

The scaling of multi-core processors poses a challenge to memory system design. While process technology scaling provides greater numbers of cores on each chip, the transistor performance and power gains that traditionally accompanied process scaling have largely ceased. Scaling performance with increased core counts must be achieved under the same or reduced energy and power budgets. Furthermore, increased cores generate more accesses to shared caches causing conflict misses as unrelated processes compete for the same cache sets. Each miss represents significant waste: wasted time as the requested data is transferred from a slow main memory, wasted energy and bandwidth when transferring cache block words that will ultimately go unused. In this talk I will explore the means to leverage memory reference speculation to reduce waste and improve efficiency in multi-core processor memory systems. In particular, I will show how control flow and effective address speculation can be used in a novel, highly accurate, prefetch engine that improves IPC by 31% versus baseline for multiprogrammed workloads, outperforming the best competing design of class with 65% less hardware state overhead. I will also demonstrate a technique to speculate on word usage within processor cache lines to eliminate 36% of the dynamic energy in the multi-core processor interconnect with no performance impact. Finally, I will discuss some on-going work in path-speculation based prefetching for the last-level cache.

Biography

Paul V. Gratz is an Associate Professor in the department of Electrical and Computer Engineering at Texas A&M University. His research interests include energy efficient and reliable design in the context of high performance computer architecture, processor memory systems and on-chip interconnection networks. He received his B.S. and M.S. degrees in Electrical Engineering from The University of Florida in 1994 and 1997 respectively. From 1997 to 2002 he was a design engineer with Intel Corporation. He received his Ph.D. degree in Electrical and Computer Engineering from the University of Texas at Austin in 2008. His papers "A Control-Theoretic Approach for Energy Efficient CPU-GPU Subsystem in Mobile Platforms" and "B-Fetch: Branch Prediction Directed Prefetching for Chip-Multiprocessors" were nominated for best papers at DAC '15 and MICRO '14 respectively. At ASPLOS '09, Dr. Gratz received a best paper award for "An Evaluation of the TRIPS Computer System." In Spring 2010, he received the "Teaching Excellence Award - Top 5%" award from the Texas A&M University System for a graduate-level computer architecture course he developed.

For more information please contact: Houman Homayoun, hhomayou@gmu.edu