Electrical & Computer ENGINEERING **SEMINAR**

Thursday, February 5
Scaife Hall Auditorium
Room 125
4:30 p.m.
Refreshments at 4:00 p.m.

Onur Mutlu
Assistant Professor
CMU ECE

Onur Mutlu is an Assistant Professor of ECE at Carnegie Mellon University. He is interested in computer architecture and systems research, especially in the interactions between languages, operating systems, compilers, and microarchitecture. He has a PhD and an MS in ECE from the University of Texas at Austin and BS degrees in Computer Engineering and Psychology from the University of Michigan, Ann Arbor.

Prior to **Carnegie Mellon**, Onur worked at **Microsoft Research** (from 2006 to 2009), **Intel Corporation**, and Advanced Micro Devices. He was a recipient of the Intel PhD fellowship in 2004, the University of Texas George H. Mitchell Award for Excellence in Graduate Research in 2005, the Microsoft Gold Star Award in 2008, and five "computer architecture top pick" paper awards by the IEEE Micro magazine.

Preventing Memory Performance Attacks in Multi-Core System

Current multi-core processors are vulnerable to a new class of denial of service attacks because of their shared memory systems. A low-importance application can intentionally or unintentionally destroy the performance of a high-importance one running on the same chip. We call such an aggressive application a memory performance hog (MPH). With the widespread deployment of multi-core systems in commodity computers and with the rapid increase in the number of on-chip cores, we expect MPHs to become a prevalent security and performance issue that could affect almost all computer users.

I will describe our solution to the problem, **"stall-time fair memory scheduling (STFM),"** that provides performance-fairness to threads sharing the memory system. STFM's goal is to equalize the memory-related slowdown experienced by equal-priority threads due to interference from other threads, while also improving overall system performance. STFM not only effectively contains memory performance hogs and provides fair memory access but also improves system utilization. I will describe how STFM is seamlessly exposed to the system software to provide different levels of memory service to threads with different priorities/weights and to enable the enforcement of different software-based Quality-of-Service policies. Our evaluations show that STFM provides a flexible, low-cost, and high-performance fairness substrate for multi-core memory systems. The final part of this talk will cover improvements over STFM and a brief discussion of our future directions on building secure and controllable multi-core systems.

ECE Seminar Hosts

Yi Luo	y1827@andrew.cmu.edu
Marios Savvides	marioss@andrew.cmu.edu
Bruno Sinopoli	brunos@andrew.cmu.edu

Carnegie Mellon

Electrical & Computer ENGINEERING