HARVEY MUDD COLLEGE

Engineering Seminar Program
Wednesday, April 13, 2011
Galileo McAlister 4:15 p.m.

Matthew Watkins, Visiting Professor

Matthew Watkins is currently a visiting assistant professor in the Engineering Department at Harvey Mudd College and will be joining Intel's architecture design team this summer. His interests are in the area of computer architecture and digital design with a special focus on reconfigurable architectures for chip multiprocessors and the design of large and heterogeneous many-core systems. He received his PhD and MS degrees in electrical and computer engineering from Cornell University and BS degrees in Computer Engineering and Electrical Engineering from the University at Buffalo. He was a 2005 National Science Foundation Graduate Research Fellow. He's a member of IEEE and the ACM.

Reconfigurable Architectures for Chip Multiprocessors

The microprocessor industry has undergone a radical change in the last 5 years. Instead of improving the performance of a single core, manufactures now incorporate multiple cores on a single die in an attempt to continue Moore's Law style performance growth. The best organization for future large-scale multicore chips remains an open question. One option that is gaining increased interest is to incorporate reconfigurable hardware on-chip to allow dynamic creation of hardware specifically tailored to accelerate the running applications.

Prior research in reconfigurable computing involved augmenting a single processor core with reconfigurable logic. Despite significant performance gains for some applications, the area and power costs can easily outweigh the benefits, especially for workloads that do not make good use of the fabric. Moreover, this prior work almost exclusively focused on uniprocessor systems and did not address the unique requirements of parallel applications.

This talk proposes a novel reconfigurable architecture for chip multiprocessors (CMPs) running a mix of serial and parallel applications. In our approach, the reconfigurable fabric is shared among multiple threads to amortize the area and power costs and increase fabric utilization. To further reduce the overhead, we propose a heterogeneous CMP with different regions optimized for different tasks, including regions with shared reconfigurable fabrics, and other regions with only conventional cores. Within a reconfigurable region, the architecture dynamically manages the assignment of threads to fabric clusters and the dynamic partitioning of the fabric and includes mechanisms that accelerate parallel applications and enable parallelization of otherwise sequential applications.