

Thursday, February 25th

Scaife Hall Auditorium

Room 125

4:30 p.m.

Refreshments at 4:00 p.m.



Joel S. Emer

Intel Fellow, Intel Architecture Group
Director, Microarchitecture Research
INTEL CORPORATION

Joel Emer is an Intel Fellow and Director of Microarchitecture Research at Intel in Hudson, Massachusetts. Previously he worked at Compaq and Digital Equipment Corporation where he held various research and advanced development positions investigating processor micro-architecture for a variety of VAX and Alpha processors and developing performance modeling and evaluation techniques. His research included pioneering efforts in simultaneous multithreading and early contributions on the now pervasive quantitative approach to processor evaluation. His current research interests include memory hierarchy design, processor reliability, reconfigurable logic-based computation and performance modeling. In his spare time, he serves as visiting faculty at MIT. He received his PhD in electrical engineering from the University of Illinois at Urbana-Champaign in 1979. He is a Fellow of both the ACM and the IEEE, and was the 2009 recipient of the Eckert-Mauchly award for lifetime contributions in computer architecture.

An Evolution of General Purpose Processing: Reconfigurable Logic Computing

The historical improvements in the performance of general-purpose processors have long provided opportunities for application innovation. Word processing, spreadsheets, desktop publishing, networking and various game genres are just some of the many applications that have arisen because of the increasing capabilities and the versatility of general-purpose processors. Key to these innovations is the fact that general-purpose processors do not predefine the applications that they are going to run. Currently, the capabilities of individual general-purpose processors are encountering challenges, such as diminishing returns in exploiting instruction-level parallelism and power limits. Thus uniprocessor performance is not scaling as rapidly as it had been. As a consequence, a variety of approaches are being employed to address this situation, including vectors, multi-threading, multi-processors, and dedicated accelerators. While each of these techniques extends a processor's compute capabilities each sacrifices generality in one way or another.

At the far extreme of non-generality is dedicated logic, which when applied to a specific task will invariably out perform and out power/perform a general-purpose processor. However, the time, difficulty and cost of special purpose design preclude dedicated logic from serving as a viable avenue for application innovation. Recently, a middle ground between fully general-purpose computing and dedicated logic has been showing increasing promise. In specific, reconfigurable logic, typically in the form of FPGAs, addresses many of the cost-related liabilities of dedicated logic and is increasingly being applied to general computation problems. In this talk, we will examine the possibilities for reconfigurable logic as a foundation for general-purpose computation. We will look at its potential and attempt to provide an analogy between the state of reconfigurable logic computing today and the early days of conventional computing. In that light, we consider how reconfigurable logic might recapitulate the history of general-purpose computation by looking at how it can fit into the architectural framework that we have generally reserved for conventional processors, how it can be more seamlessly be incorporated into a system and how reconfigurable logic might be made more efficient and be more effectively programmed by looking at the semantic gap between programming languages and the compute fabric itself.

ECE Seminar Hosts

Jeyanandh Paramesh paramesh@ece.cmu.edu

Onur Mutlu onur@cmu.edu

Gabriela Hug ghug@ece.cmu.edu