

# *Dynamically Heterogeneous Cores Through 3D Resource Pooling*



**SF STATE**

*Houman  
Homayoun*  
NSF Computing  
Innovation Fellow  
*UC San Diego*



**Feb 13, 2012  
Monday**

**1.10-2pm**

**SCI 256**

## *Synopsis:*

3D die stacking is a recent technological development which makes it possible to create chip multiprocessors using multiple layers of active silicon bonded with low-latency, high-bandwidth, and very dense vertical interconnects. 3D die stacking technology provides very fast communication, as low as a few picoseconds, between processing elements residing on different layers of the chip. The rapid communication network in a 3D stack design, along with the expanded geometry, provides an opportunity to dynamically share on-chip resources among different cores. This research describes an architecture for a dynamically heterogeneous processor architecture leveraging 3D stacking technology. Unlike prior work in the 2D plane, the extra dimension makes it possible to share resources at a fine granularity between vertically stacked cores. As a result, each core can grow or shrink resources, as needed by the code running on the core. This architecture, therefore, enables runtime customization of cores at a fine granularity and enables efficient execution at both high and low levels of thread parallelism. This architecture achieves performance gains of up to 2X, depending on the number of executing threads, and gains significant advantage in energy efficiency.

## *Speaker Bio:*

**Houman Homayoun** is currently a 2010 NSF Computing Innovation (CI) Fellow awarded by the Computing Research Association (CRA) and the Computing Community Consortium (CCC). His research is on power-temperature and reliability-aware memory and processor design optimizations and spans the areas of computer architecture, embedded systems, circuit design, and VLSI-CAD, where he has published more than 30 technical papers on the subject, including some of the earliest work in the field to address the importance of cross-layer power and temperature optimization in memory peripheral circuits. He currently works at the University of California, San Diego leading a number of research projects funded by the NSF, SRC, Intel, and Google, including the design of the next generation 3D heterogeneous multicores, low power hybrid SRAM-NVM memory hierarchy design, reliability-aware cache design, and peak power management in data centers. Houman was a recipient of the four-year UC-Irvine Computer Science Department chair fellowship. He received his PhD degree from the Department of Computer Science at the University of California, Irvine in 2010, an MS degree in computer engineering in 2005 from University of Victoria, Canada and his BS degree in electrical engineering in 2003 from Sharif University of technology, Tehran, Iran..