Title: Porting Performance across GPUs and FPGAs through Multilevel Granularity Synthesis

ABSTRACT:
As growing power dissipation and thermal effects disrupted the rising clock frequency trend and threatened to annul Moore's law, the computing industry has switched its route to higher performance through parallel processing. GPUs and FPGAs are becoming popular systems for speeding up computation-intensive kernels of scientific, imaging, and simulation applications. GPUs can execute thousands of concurrent threads, while FPGAs provide customized concurrency for highly parallel kernels. However, exploiting the parallelism available in these applications is currently not a push-button task. Often the programmer has to expose the application's fine- and coarse-grained parallelism by using special programming languages. CUDA and OpenCL are such parallel languages that are driven by the GPU industry and are gaining popularity. Based on a collaborative effort from UIUC, UCLA and ADSC, we have been working on a new CUDA-to-FPGA design flow called FCUDA, which efficiently maps parallelism exposed in CUDA onto the reconfigurable fabric. Such a design flow would enable high parallelism expression, common frontend language for heterogeneous compute platforms, and performance porting between GPU and FPGA. Experimental results show that FCUDA can transform GPU software kernels to FPGA implementations that offer competitive performance compared to GPU executions but at a fraction of the energy cost.

BIO:
Dr. Deming Chen obtained his BS in computer science from the University of Pittsburgh, Pennsylvania in 1995, and his MS and PhD in computer science from the University of California at Los Angeles in 2001 and 2005 respectively. He worked as a software engineer for 1995-1999 and 2001-2002. He is an associate professor in the ECE department of the University of Illinois at Urbana-Champaign. His current research interests include high-level synthesis, nano-systems design and nano-centric CAD techniques, reconfigurable computing, GPU computing, and microarchitecture and SoC design under parameter variations. He is a technical committee member for a series of conferences and a committee or CAD track chair for several conferences. He is an associate editor for TCAD, TVLSI, TCAS-I, JCSC, and JOLPE. He obtained the Achievement Award for Excellent Teamwork from Aplus Design Technologies in 2001, the Arnold O. Beckman Research Award from UIUC in 2007, the NSF CAREER Award in 2008, and four Best Paper Awards from various conferences. He received the ACM SIGDA Outstanding New Faculty Award in 2010. He is included in the List of Teachers Ranked as Excellent in 2008. He is a senior member of IEEE.