

Accelerating System-Level Design Tasks using Graphics Processors (Half Day)

Udeeptha D. Bordoloi, *Advanced Micro Devices (AMD), Sunnyvale, USA*
Samarjit Chakraborty, *TU Munich, Germany*
Unmesh D. Bordoloi, *Linköping University, Sweden*

Abstract

Recent years have seen the increasing use of graphics processing units (GPUs) for non-graphics related applications. Applications that have harnessed the computational power of GPUs span across numerical algorithms, computational geometry, database processing, image processing, astrophysics and bioinformatics. There are many compelling reasons behind exploiting GPUs for such general-purpose computing tasks. First, modern GPUs are extremely powerful. For example, high-end GPUs such as the NVidia GeForce GTX 480 and ATI Radeon 5870 have 1.35 TFlops and 2.72 TFlops of peak single precision performance, whereas a high-end general-purpose processor such as the Intel Core i7-960 has a peak performance of 102 Gflops. Additionally, the memory bandwidth of these GPUs are more than 5x greater than what is available to a CPU, which allows them to excel even in low compute intensity but high bandwidth usage scenarios. Second, GPUs are now commodity items as their costs have dramatically reduced over the last few years.

In spite of a wide variety of computationally expensive system-level design tasks (in the context of embedded systems design) that are regularly solved by software tools running on desktops and laptops equipped with high-end GPUs, the use of GPUs for accelerating such problems is still not a conventional practice within the design automation community. As a result, of late, there has been a lot of research interest in demonstrating the applicability of GPUs in accelerating design automation tasks. Some of tasks that have been accelerated using modern GPUs include schedulability/timing analysis, hardware/software partitioning, fault simulation, and verification of digital designs. In this tutorial we will describe techniques for programming GPUs for general purpose computing (i.e., non-graphics applications) and cover a number of case studies from the electronic design automation area. We will demonstrate how GPUs can lead to significant improvement in running times and hence the usability of the design tools that exploit them. In particular, we will start by introducing the graphics processor architecture and programming models for GPUs (OpenCL and CUDA). OpenCL is an open standard for programming GPUs (and also other modern processors) and is a cross-platform alternative to CUDA. It has been created by a consortium that includes AMD, Apple, IBM, Intel, and Nvidia. We will then discuss various examples of system-level design tasks and identify their computational kernels. Finally, we will present different case studies to illustrate how system-level design algorithms have to be suitably modified in order to map them onto GPUs.

Speaker Biographies

Udeeptha D. Bordoloi is a Senior Member of Technical Staff at AMD working on GPGPU technology. He is an expert in high performance computation and GPUs, and has helped advance and optimize GPU utilization in multiple large-scale GPU based clusters in the top500 supercomputers list. He has published papers, tutorials and lectures promoting the state-of-the-art in GPUs, and collaborates with the academia and the industry on the newest features of the latest GPUs being developed at AMD. He obtained his MS degree in Electrical Engineering from Washington University in St. Louis and his PhD in Computer Science from The Ohio State University. Since then he has worked in the field of graphics and visualization with industry leading companies and also with bootstrapping startups. He has many academic publications and patents to his name and continues to collaborate with and contribute to the research community. Besides GPGPU and HPC, he has interests in Image Processing, Graphics and Visualization.

Samarjit Chakraborty is a Professor of Electrical Engineering at TU Munich in Germany, where he heads the Institute for Real-Time Computer Systems (RCS). Prior to joining TU Munich, from 2003 – 2008 he was an

Assistant Professor of Computer Science at the National University of Singapore. He obtained his PhD in Electrical and Computer Engineering from ETH Zurich in 2003. His research interests cover all aspects of system-level design of real-time embedded systems and software, and his work has attracted funding from companies like Intel, General Motors, Bosch and BMW. Prof. Chakraborty regularly serves on the technical program committees of several conferences in the area of real-time and embedded systems, such as DATE, CODES+ISSS, ASP-DAC, RTSS and RTAS. He was the TPC Co-Chair of EMSOFT 2009, and is currently serving as the General Co-Chair of Embedded Systems Week 2010. Apart from more than 80 referred journal and conference articles, he has authored a number of patents and regularly gives invited talks and tutorials at various research labs and international conferences (such as ESWeek, VLSI Design, ACM Multimedia and ICME). For his Ph.D. thesis, he received the ETH Medal and the European Design and Automation Association's "Outstanding Doctoral Dissertation Award" in 2004. His work has also received a HiPEAC Paper Award in 2009 and Best Paper Award nominations at DAC 2005, CODES+ISSS 2006, ECRTS 2007, and CODES+ISSS 2008.

Unmesh D. Bordoloi is a post-doctoral researcher at Linköping University, Sweden. Prior to joining Linköping he spent one year as a post-doctoral researcher at the Verimag Laboratory in Grenoble, France. He obtained his PhD from the School of Computing at the National University of Singapore in 2008. He was awarded the "President's Graduate Fellowship" and the "Sun Microsystems Ph.D. Fellowship" for his research work at the National University of Singapore. He received a Bachelor of Technology in Computer Science and Engineering from the National Institute of Technology, Rourkela in 2004. In 2007, he visited General Motors R&D, India Science Lab at Bangalore, as an intern. His research focuses on system-level design and analysis issues like schedulability analysis, hardware/software co-design, and fault-tolerance. His work has appeared in a number of top-tier conferences and journals such as DAC, RTAS and the ACM Trans. Embedded Computing Systems. He is currently investigating techniques for automatically mapping a large class of algorithms from the design automation domain onto GPUs.