SEMINAR NOTICE:

An Accurate GPU Performance Model for Effective Control Flow Divergence Optimization

Yun (Eric) Liang

Research Scientist at the Advanced Digital Sciences Center in Singapore

Thursday, May 31
3:00 p.m.
CSL 301

Abstract:

Graphics processing units (GPUs) are increasingly critical for general-purpose parallel processing performance. GPU hardware is composed of many streaming multiprocessors, each of which employs the single-instruction multiple-data (SIMD) execution style. This massively parallel architecture allows GPUs to execute tens of thousands of threads in parallel. Thus, GPU architectures efficiently execute heavily data-parallel applications. However, due to this SIMD execution style, resource utilization and thus overall performance can be significantly affected if computation threads must take diverging control paths. Control flow divergence in GPUs is a well-known problem. However, prior approaches are seriously affected by the lack of a guiding metric that properly estimates how control flow divergence affects application performance. In this talk, I will first present a case study of accelerating 3D sound localization using GPU. Then, I will introduce a metric that simply and accurately estimates performance of computation-bound GPU kernels with control flow divergence, and use the metric as a value function for thread re-grouping algorithms.

Biosketch:

Yun (Eric) Liang is a Research Scientist with the Advanced Digital Sciences Center (ADSC). His research interests include embedded system, computer architecture and compiler. He has published 20 research papers in the top conferences on embedded systems, computer architecture, real-time systems, and hardware, such as DAC, FCCM, CASES, RTSS, IPDPS, CODES+ISSS, DATE, and FPT. His work has received the Best Paper Award of FCCM 2011 and Best Paper Award nominations from DAC 2012, FPT 2011, and CODES+ISSS 2008. Currently, Eric is working on a project related to FPGA and GPU in ADSC, led by Professor Deming Chen in UIUC.