



I2PC

Illinois-
Intel
Parallelism
Center



Distinguished Speaker Series

Matt Pharr

Intel Corporation

ispc: A High-Performance SPMD Compiler for The CPU

Thursday, March 15, 2012

4 - 5 PM CST

3405 Siebel Center

For more information about I2PC
research, education, and activities, visit:
i2pc.cs.illinois.edu

Attend this talk remotely at:
<http://media.cs.illinois.edu/live/I2PClive.aspx>

A chat interface has been set up for questions:
<http://i2pc.cs.illinois.edu/chat>

Abstract

SIMD parallelism has become an increasingly important mechanism for delivering performance in modern CPUs, due its power efficiency and relatively low cost in die area compared to other forms of parallelism. Unfortunately, languages and compilers for CPUs haven't kept up with the hardware's capabilities. Existing CPU parallel programming models have almost entirely focused on multicore parallelism, neglecting the substantial computational capabilities available in SIMD vector units, while GPU-oriented languages like OpenCL suffer from constraints that impair ease of use on CPUs and lack capabilities needed to achieve maximum efficiency on CPUs due to their focus on GPU architectures.

This talk describes a new compiler, Intel SPMD Program Compiler (ispc), that delivers very high performance on CPUs thanks to effective use of both processor multiple cores and SIMD vector units. ispc draws from GPU programming languages, which have shown that for many applications, the easiest way to program SIMD units is to use a single-program, multiple-data (SPMD) model, with one instance of the program mapped to each SIMD lane. This talk will describe discuss language features that make ispc easy to adopt and use productively with existing software systems and present results showing that ispc delivers up to 35x speedups on a 4-core system and up to 240x speedups on a 40-core system for complex workloads compared to serial C++ code.

Bio

Matt Pharr is a Principal Engineer at Intel where he most recently has built ispc, the Intel SPMD Program Compiler. He previously was the lead lead graphics architect in the Advanced Rendering Technology group, which researched new interactive graphics algorithms and programming models. Prior to coming to Intel, he co-founded Neoptica, which developed new programming models for graphics on heterogeneous CPU+GPU systems; Neoptica was acquired by Intel.

Before Neoptica, Matt was in the Software Architecture group at NVIDIA, co-founded Exluna (acquired by NVIDIA), worked in Pixar's Rendering R&D group, and received his Ph.D. from the Stanford Graphics Lab. With Greg Humphreys, he wrote the textbook "Physically Based Rendering: From Theory to Implementation", which has been used for graduate-level graphics classes at over 20 universities.



ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN