Illinois-
Intel
Parallelism
Center

I2PC

Distinguished Speaker Series

Zhiyuan Li

Purdue University

How to Make Loop-level Data Dependence Profiling Fast and Memory Efficient?

Thursday, February 23, 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.asx

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

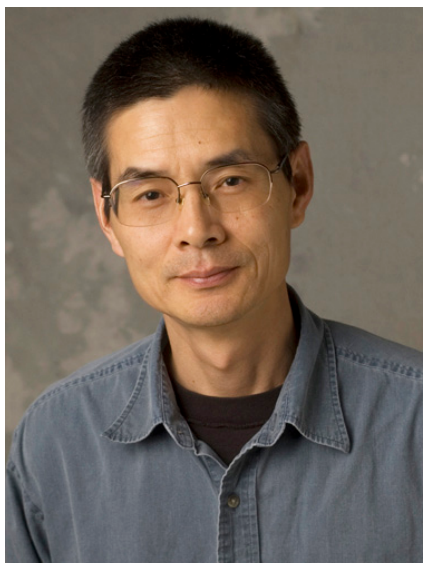

Abstract

Execution-driven data dependence profiling for nonnumerical programs has gained significant interest in recent years because it can resolve memory access ambiguity exactly through program execution, which allows data dependences to be analyzed exactly in spite of complicated aliases and referencing expressions which often defeat the compiler. Dependence result obtained through profiling is exact for specific inputs only. Nonetheless, it provides valuable insights both to compiler designers (who may discover ways to improve their compilers) and programmers (who may discover ways to parallelize the program by hand).

Unfortunately, dependence profiling itself can take tremendous memory and machine time. For practical use, both the memory efficiency and the execution speed need to be improved by at least one or two orders of magnitude from the state of the art. We explore ways to make such improvements with the help of several compiler and runtime techniques. These methods include a) parallelized profiling (with various granularities), b) analysis of type consistency and aliasing to allow a method to embed memory tag with the original data structure instead of using the conventional hash table, c) a partial dependence graph that is proven to be sufficient for loop transformation and parallelization. Experimental data for SPEC CPU 2006 benchmarks are presented.

Bio

Dr. Zhiyuan Li joined Purdue University in 1997 where he is now a professor in the Department of Computer Science. He received his PhD degree in Computer Science in 1989 from University of Illinois, Urbana-Champaign where he returned to work in 1990-1991 as a senior software engineer in Center for Supercomputing Research and Development after one year teaching at York University, Canada. From 1991 to 1997 he taught at Department of Computer Science, University of Minnesota. Dr. Li currently takes on several research projects sponsored by NSF and by Intel, ranging from petascale numerical applications, compiler tools for multicore machines, compiler tools for reliable networked embedded systems, and solutions for memory bottleneck problems on multicore.

ILLINOIS
UNIVERSITY OF ILLINOIS AT URBANA-CHAMPAIGN