

Compilation and Optimization for High Performance Computing

Kleanthis Psarris
Department of Computer Science
The University of Texas at San Antonio
San Antonio, TX 78249
psarris@cs.utsa.edu

Abstract

High end parallel and multi-core processors rely on compilers to perform the necessary optimizations and exploit concurrency in order to achieve higher performance. However, source code for high performance computers is extremely complex to analyze and optimize. In particular, program analysis techniques often do not take into account complex expressions during the data dependence analysis phase. Most data dependence tests are only able to analyze linear expressions, even though non-linear expressions occur very often in practice. Therefore, considerable amounts of potential parallelism remain unexploited. In this talk we propose new data dependence analysis techniques to handle such complex instances of the dependence problem and increase program parallelization. Our method is based on a set of polynomial time techniques that can prove or disprove dependences in source codes with non-linear and symbolic expressions, complex loop bounds, arrays with coupled subscripts, and if-statement constraints. In addition our algorithm can produce accurate and complete direction vector information, enabling the compiler to apply further transformations. To validate our method we performed an experimental evaluation and comparison against the I-Test, the Omega test and the Range test in the Perfect and SPEC benchmarks. The experimental results indicate that our dependence analysis tool is accurate, efficient and more effective in program parallelization than the other dependence tests. The improved parallelization results into higher speedups and better program execution performance in several benchmarks.