Optimal Superblock Scheduling Using Enumeration

Ghassan Shobaki
Computer Science Department
University of California, Davis
shobaki@cs.ucdavis.edu

Kent Wilken
Electrical and Computer Engineering Department
University of California, Davis
wilken@ece.ucdavis.edu

Abstract

The superblock is a scheduling region that is used by compilers for exploiting instruction-level parallelism across basic blocks. Many heuristic techniques have been proposed for solving this difficult scheduling problem, but none accurately approximates the optimal solution. This paper presents a new technique that finds provably optimal solutions to superblock scheduling problems. The technique is based on reducing the problem of finding branch combinations that yield incrementally increasing weighted execution times to a subset-sum problem, which is solved by dynamic programming. An enumerative approach that employs a number of powerful pruning techniques to efficiently explore the solution space is then used to search for a feasible schedule for each branch combination. Experimental evaluation using the SPEC CPU fp2000 and int2000 benchmarks shows that, within a per-problem time limit of one second, this combination of dynamic programming and enumeration optimally solves about 99% of the hard superblock scheduling problems with an average solution time of 9 milliseconds per problem. For 80% of the hard problems, the optimal schedule is improved compared to the schedule produced by an established heuristic technique.

Keywords: global instruction scheduling, compiler optimizations, superblock, optimal scheduling, enumeration.