

Compiler-Assisted Cache Replacement: Problem Formulation and Performance Evaluation

Hongbo Yang, R. Govindarajan, Guang R. Gao and Ziang Hu

To appear at the *16th Workshop on Languages and Compilers for Parallel Computing (LCPC03)*, College Station, TX, 2-4 October 2003

Abstract

Recent research results show that conventional hardware-only cache solutions result in unsatisfactory cache utilization for both regular and irregular applications. To overcome this problem, a number of architectures introduce instruction hints to assist cache replacement. For example, Intel Itanium architecture augments memory accessing instructions with cache hints to distinguish data that will be referenced in the near future from the rest. With the availability of such methods, the performance of the underlying cache architecture critically depends on the ability of the compiler to generate code with appropriate cache hints. In this paper we formulate this problem – giving cache hints to memory instructions such that cache miss rate is minimized – as a 0/1 knapsack problem, which can be efficiently solved using a dynamic programming algorithm. The proposed approach has been implemented in our compiler testbed and evaluated on a set of scientific computing benchmarks. Initial results show that our approach is effective on reducing the cache miss rate and improving program performance.