Efficient Design Space Exploration of High Performance Embedded Out-of-Order Processors

Stijn Eyerman Lieven Eeckhout Koen De Bosschere
ELIS, Ghent University, Sint-Pietersnieuwstraat 41, B-9000 Gent, Belgium
{seyerman,leeckhou,kdb}@elis.UGent.be

Abstract

Previous work on efficient customized processor design primarily focused on in-order architectures. However, with the recent introduction of out-of-order processors for high-end high-performance embedded applications, researchers and designers need to address how to automate the design process of customized out-of-order processors. Because of the parallel execution of independent instructions in out-of-order processors, in-order processor design methodologies which subdivide the search space in independent components are unlikely to be effective in terms of accuracy for designing out-of-order processors. In this paper we propose and evaluate various automated single- and multi-objective optimizations for exploring out-of-order processor designs. We conclude that the newly proposed genetic local search algorithm outperforms all other search algorithms in terms of accuracy. In addition, we propose two-phase simulation in which the first phase explores the design space through statistical simulation; a region of interest is then simulated through detailed simulation in the second phase. We show that simulation time speedups can be obtained of a factor 2.2X to 7.3X using two-phase simulation.

1. Introduction

High performance embedded applications such as multimedia, networking, imaging, high end consumer applications, etc. are an important market segment today. And because of the high performance requirements of these embedded applications, computer companies move into (or did already move into) using out-of-order processors. Examples are NEC's VR55000 and VR77100 Star Sapphire and SandCraft's SR710X0 64-bit MIPS processors. For these high end embedded systems it is extremely important to optimize the design for a given application or a given set of applications with a limited engineering effort. The design question then is how to organize the microprocessor's microarchitecture, *i.e.*, how big should the caches be, what should the processor width be, how many in-flight instructions should there be, etc.

A lot of prior work has been done on automated design space exploration methodologies for in-order (incl. VLIW) architectures, see for example PICO [8], Sherpa [13] and the references therein. The in-order processor design methodologies often assume that the whole system can be broken up into independent subcomponents. Each of those subcomponents then gets optimized before reassembling the whole system. Because of the complex interactions in out-of-order processors such as parallel executions of independent instructions, latency hiding mechanisms, speculative execution, etc., it is unlikely that these design methodologies can be used for customized out-of-order processor design.

In this paper we propose the use of automated design space exploration for efficiently exploring the out-of-order processor design space. We evaluate various single-objective and multi-objective search algorithms and propose genetic local search as a new single-objective search algorithm that outperforms all other search algorithms in terms of accuracy. In addition, we also propose two-phase simulation to prune the design space. The idea is to explore the design space quickly using a fast simulation technique as a first step. And we propose statistical simulation for this purpose. We then explore a small region of interest identified in the first phase through detailed (and thus slower) cycle-accurate architectural simulation. The overall simulation speedup obtained through two-phase simulation is 2.2X to 7.3X in our experiments. And we show that two-phase simulation can be applied to single-objective as well as to multi-objective optimization criteria.

2. Previous work

There exists a lot of work on design space exploration. Obviously, we cannot discuss the whole literature written on this topic. For a detailed description on design space exploration techniques, we refer to [5]. Since our main contributions are on design space pruning and automated design space exploration, we will revisit some of those.

2.1. Design space pruning

A first approach to reducing the total simulation time is to prune the design space through what could be called hierarchical design space exploration. Hekstra *et al.* [6] for example describe how they explored the TriMedia CPU64 design. In order to explore the huge design space, they first probe the design space in order to identify the architectural parameters that affect overall performance the most. These