

An Evaluation of Autotuning Techniques for the Compiler Optimization Problems

Amir Hossein Ashouri, Gianluca Palermo and Cristina Silvano

Politecnico di Milano,
Milan, Italy

{amirhossein.ashouri, ginaluca.palermo, cristina.silvano}@polimi.it

Abstract. Diversity of today's architectures have forced programmers and compiler researchers to port their application across many different platforms. Compiler auto-tuning itself plays a major role within that process as it has certain levels of complexities that the standard optimization levels fail to bring the best results due to their average performance output. To address the problem, different optimization techniques has been used for traversing, pruning the huge space, adaptability and portability. In this paper, we evaluate our different autotuning approaches including the use of Design Space Exploration (DSE) techniques and machine learning to further tackle the both problems of selecting and the phase-ordering of the compiler optimizations. It has been experimentally demonstrated that using these techniques have positive effects on the performance metrics of the applications under analysis and can bring up to 60% performance improvement with respect to standard optimization levels (e.g. -O2 and -O3) on the selection problem and up to 4% w.r.t. to LLVM's standard optimization on the phase-ordering problem.

1 Introduction

Conventional software applications are first developed in the desired high-level source-code (e.g. C, C++) and then are passed through the compilation phase to build the executable. The later phase includes compiler optimization process in which the target metrics such as execution time, code-size, power, etc are optimized depending on the desired scenario. Compiler optimizations are playing an important role to transform the source-code to an optimized variation. Usually, open-source/industrial compiler platforms are coming off-the-shelf with some standard optimization levels (e.g. -O1, -O2, -O3 or -Os) to bring the average-good results for conventional platforms. However, quite often they fail to bring the optimal results for specific applications, architectures and platforms. In the short paper, two different techniques for compiler auto-tuning, namely, DSE and Machine Learning based techniques have been proposed to accommodate and address the problem of selecting the best compiler optimization for a given application.

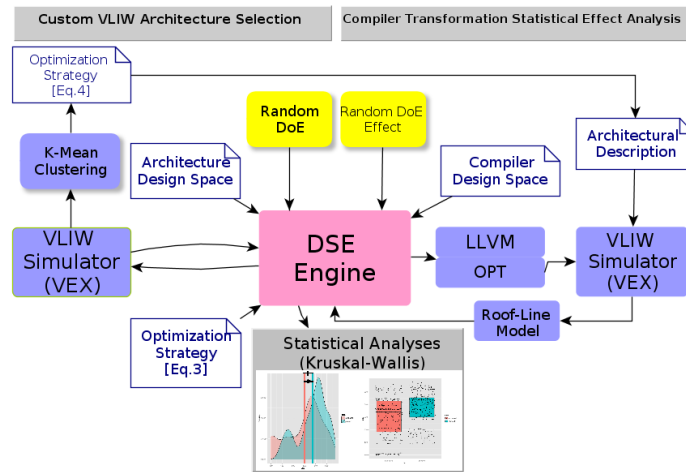


Fig. 1: Approach (I): DSE Proposed Methodology

2 Approach (I): DSE Approach

Design Space Exploration (DSE) refers to the activity of exploring the design parameters alternatives before the actual design. It deals with pruning and exploring the design space efficiently. The proposed work targets the exploration of compiler options parameters, in order to automatically explore the design space and analyze the compiler-architecture co-design. We experimentally assessed the proposed methodology in Very-long-Instruction-Word (VLIW) architecture by applying random Design of Experiment (DoE) and an automatic tool-chain including our Multi-Objective System Tuner tool (MOST), a wrapper and a compiler/simulator; namely, LLVM and VLIW-EXample (VEX). It enables to automatically explore, optimize and analyze the optimizations by using several standard benchmarks for both high-end embedded and signal processing applications [1]. Analytically, we show that the adoption of the specific methodology either in a cross-architecture and/or cross-application manner, can deliver significant application specific insights thus enabling the designer to guide through decisions regarding the architecture and the compilation optimization strategy [2]. Figure 1 represents the proposed methodology for compiler co-exploration with DSE. The work-flow starts by inferring the *Pareto-optimal* architectural design space and then feeding the found architectural properties to the compiler framework. Statistical analyses will be applied at the end to assess the correlation between utilizing the certain compiler options and the observed performance metrics. Figure 2 is showing different distributions derived by applying the proposed DSE.




Fig. 2: Visualization of (a) *licm*'s significant positive effect, (b) *reassociate*'s no significant effect

3 Approach (II): Machine Learning Approach

Diversity in applications and architecture, simply makes it barely possible to manually optimize and port the source-codes for each application/architecture. *Random Iterative Compilation* fails to efficiently bring the optimal results due to its high demand on time and number of iterations. In order to improve the portability of compiler optimization with respect to the handcrafted approaches, machine learning has been used to address both the *selection of compiler optimization options* and *phase-ordering problem* [3] to predict the right optimization to be applied given an unseen application [4].

3.1 The Problem of Selecting the Right Set of Compiler Optimizations

Addressing the issue on the second approach, we propose a Machine Learning based autotuning framework that maximizes the performance of a target application. COBAYN: Compiler Autotuning Framework Using Bayesian Networks, starts by applying statistical methodology with Bayesian Networks to infer the probability distribution of the compiler optimizations to be enabled to achieve the best performance. We start to drive the iterative compilation process by sampling from the probability distribution. Likewise most machine learning approaches, here we use a couple of sets of training applications to learn the statistical relations between application features and the compiler optimizations. Given a new unseen application, its features are fed into the machine learning algorithm as *evidence* on the distribution. This evidence imposes a bias on the distribution. Since compiler optimizations are correlated with the software features, we can redo the process of sampling for the new target application. Figure 3 demonstrates the second proposed approach that is assessed on an embedded *ARM* device with GCC compiler. The obtained probability distribution is indeed application-specific and effectively exploits the use of iterative compilation process as it only drives with the most promising compiler optimizations [5, 6]. Figure 4, represents the result of our proposed methodology, COBAYN, against GCC's standard optimization levels `-O2` and `-O3` when using *cBench* suite. It represents significant speedup factor over the the experimentally tested applications with the average 56% and 47% improvement, respectively against `-O2` and

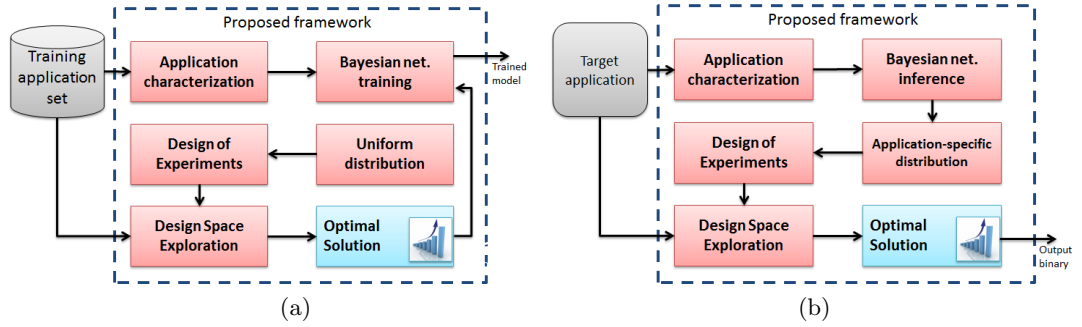


Fig. 3: Approach (II): Overview of the proposed M.L Methodology a) training phase b) predicting phase




Fig. 4: Performance improvement of our Bayesian Networks w.r.t -O2 and -O3

-O3. COBAYN, led to reach a factor of $3\times$ exploration speedup compared with the *random iterative compilation* having a fixed number of predictions.

3.2 The Phase-ordering Problem

when taking into considerations the order of the appearance of the compiler optimizations, the so-called *phase-ordering* problem comes into play. The space gets enormously bigger and simple classic supervised techniques are not able to tackle accurate models for prediction. Addressing the phase-ordering problem, we propose an *intermediate speedup* predictor that is able to predict the current optimization to be applied given the state of the code being optimized. We used predictive models and dynamic software characterization to construct the application specific models in cross-validation manner. In order to speedup the exploration on the space, we defined two traversing heuristics, depicted in Figure 5, that use Depth First Search (DFS) and exhaustive search within the prediction space. The proposed approach reaches up to 4% speedup w.r.t LLVM’s default compilation performance [3].




Fig.5: Performance improvement of our proposed speedup prediction models against LLVM

4 Conclusion

This paper presents two main approaches for the compiler autotuning problem using DSE and machine learning. The experimental results is shown speedup on the performance metrics while classifying the effective compiler optimizations derived by DSE approach and 40% - 60% speedup with against GCC's -O2 and -O3 on an ARM embedded-board using COBAYN.

References

1. A. H. Ashouri, "Design space exploration methodology for compiler parameters in vliw processors," 2012, <http://hdl.handle.net/10589/72083>.
2. A. H. Ashouri, V. Zaccaria, S. Xydis, G. Palermo, and C. Silvano, "A framework for compiler level statistical analysis over customized vliw architecture," in *Very Large Scale Integration (VLSI-SoC), 2013 IFIP/IEEE 21st International Conference on*. IEEE, 2013, pp. 124–129.
3. A. H. Ashouri, A. Bignoli, G. Palermo, and C. Silvano, "Predictive modeling methodology for compiler phase-ordering," in *Proceedings of 7th Workshop on Parallel Programming and Run-Time Management Techniques for Many-core Architectures and 5th Workshop on Design Tools and Architectures for Multicore Embedded Computing Platforms*. ACM, 2016.
4. F. Agakov, E. Bonilla, J. Cavazos, B. Franke, G. Fursin, M. F. O'Boyle, J. Thomson, M. Toussaint, and C. K. Williams, "Using machine learning to focus iterative optimization," in *Proceedings of the International Symposium on Code Generation and Optimization*. IEEE Computer Society, 2006, pp. 295–305.
5. A. H. Ashouri, G. Mariani, G. Palermo, E. Park, J. Cavazos, and C. Silvano, "Cobayn: Compiler autotuning framework using bayesian networks," *ACM Transactions on Architecture and Code Optimization (TACO)*, vol. 13, no. 2, p. 21, 2016.
6. A. H. Ashouri, G. Mariani, G. Palermo, and C. Silvano, "A bayesian network approach for compiler auto-tuning for embedded processors," in *Embedded Systems for Real-time Multimedia (ESTIMedia), 2014 IEEE 12th Symposium on*. IEEE, 2014, pp. 90–97.