

# UC Irvine

## UC Irvine Previously Published Works

### Title

Towards an Achievable Performance for the Loop Nests

### Permalink

<https://escholarship.org/uc/item/2ch6v2p2>

### Authors

Shivam, Aniket  
Watkinson, Neftali  
Nicolau, Alexandru  
[et al.](#)

### Publication Date

2019

### DOI

10.1007/978-3-030-34627-0\_6

### Copyright Information

This work is made available under the terms of a Creative Commons Attribution License, available at <https://creativecommons.org/licenses/by/4.0/>

Peer reviewed

# Towards an Achievable Performance for the Loop Nests

Aniket Shivam<sup>1</sup>, Neftali Watkinson<sup>1</sup>, Alexandru Nicolau<sup>1</sup>,  
David Padua<sup>2</sup>, and Alexander V. Veidenbaum<sup>1</sup>

<sup>1</sup> Department of Computer Science, University of California, Irvine  
{aniketsh,watkinso,nicolau,alexv}@ics.uci.edu

<sup>2</sup> Department of Computer Science, University of Illinois at Urbana-Champaign  
padua@illinois.edu

**Abstract.** Numerous code optimization techniques, including loop nest optimizations, have been developed over the last four decades. Loop optimization techniques transform loop nests to improve the performance of the code on a target architecture, including exposing parallelism. Finding and evaluating an optimal, semantic-preserving sequence of transformations is a complex problem. The sequence is guided using heuristics and/or analytical models and there is no way of knowing how close it gets to optimal performance or if there is any headroom for improvement.

This paper makes two contributions. First, it uses a comparative analysis of loop optimizations/transformations across multiple compilers to determine how much headroom may exist for each compiler. And second, it presents an approach to characterize the loop nests based on their hardware performance counter values and a Machine Learning approach that predicts which compiler will generate the fastest code for a loop nest. The prediction is made for both auto-vectorized, serial compilation and for auto-parallelization. The results show that the headroom for state-of-the-art compilers ranges from 1.10x to 1.42x for the serial code and from 1.30x to 1.71x for the auto-parallelized code. These results are based on the Machine Learning predictions.

## 1 Introduction

Modern architectures have been evolving towards greater number of cores on the chip, as well as, improving the processing capabilities of individual cores. Each core in the current multi-core architectures includes the capability to process Single Instruction Multiple Data (SIMD) or Vector instructions. State-of-the-art compilers, or code optimizers, use advanced loop transformation techniques to modify the loop nests so as to take advantage of these SIMD instructions. The underlying code optimization techniques in the compilers to *auto-vectorize* the loop nests[19,1,26] require careful analysis of data dependences, memory access patterns, etc. Similarly, a serial version of the loop nest may be parallelized i.e. transformed such that loop iterations can be reordered and scheduled for parallel execution across the multiple cores. These transformations are characterized

as *auto-parallelization* techniques[18,13,15,14,16,3,7] and the end product is a multi-threaded code.

Some key transformations for optimizing loop nests[26,12] are Distribution, Fusion, Interchange, Skewing, Tiling and Unrolling. The best set of transformations for a given loop nest can be any possible sequence of these transformations with even repeating transformations. Even though the compilers may have the ability to perform important loop transformations, the built-in heuristics and analytical models that drive these optimizations to determine the order and the profitability of these transformations may lead to sub-optimal results. Evaluation studies[23,17,10] have shown that state-of-the-art compilers may miss out on opportunities to optimize the code for modern architectures. But a major challenge in developing heuristics and profitability models is predicting the behavior of a multi-core processor which has complex pipelines, multiple functional units, memory hierarchy, hardware data prefetching, etc. Parallelization of loop nests involve further challenges for the compilers, since communication costs based on the temporal and spatial data locality among iterations have an impact on the overall performance too. These heuristics and models differ between compilers which leads to different quality of the generated code for the loop nests and therefore, the performance may vary significantly. There are various compilers and domain specific loop optimizers that perform auto-vectorization and, in some cases, auto-parallelization such Intel ICC, GNU GCC, LLVM Clang, etc. By observing their relative performance one can identify relative headroom.

Embedding Machine Learning models in compilers is continuously being explored by the research community[6,23,24,9,22,5,25,2]. Most of the previous work used Machine Learning in the domain of auto-vectorization, phase-ordering and parallelism runtime settings. This work applies Machine Learning on a coarser level, in order to predict the most suited code optimizer - for serial as well as parallel code.

Previous studies have shown that hardware performance counters can successfully capture the characteristic behavior of the loop nests. In those studies, Machine Learning models either use a mix of static features (collected from source code at compile time) and dynamic features (collected from profiling)[23,24], or exclusively use dynamic features[6,25,2]. This work belongs to the second class and exclusively uses hardware performance counters collected from profiling a serial (-O1) version of a loop nest and uses these dynamic features as the input for the Machine Learning classifiers. It also shows that it is feasible to use hardware performance counters from an architecture to make predictions for similar multi-core architectures.

The focus of this work is to consider state-of-the art code optimizers and then use Machine Learning algorithms to make predictions for better, yet clearly achievable performance for the loop nests using these code optimizers. This is what defines a possible headroom. We believe that recognizing the inherent behavior of loop nests using hardware performance counters and Machine Learning algorithms will present an automated mechanism for compiler writers to identify where to focus on making improvements in order to achieve better performance.

## 2 Experimental Methodology

This section describes the candidate code optimizers and the architectures that we considered for this work and methodology for conducting the experiments.

### 2.1 Code Optimizers

In this work we considered 4 candidate code optimizers, as shown in Table 1, including Polly[11,20], a Polyhedral Model based optimizer for LLVM. 2 out of those 4 optimizers can perform auto-parallelization of the loop nests. The hardware performance counters are collected using an executable generated by `icc` with flags `-O1 -no-vec`, in order to disable all loop transformations, and disable vector code and parallel code generation.

### 2.2 Benchmarks

The first benchmark suite that we use for our experiment is Test Suite for Vectorizing Compilers (TSVC) as used by Callahan et al.[4] and Maleki et al.[17] for their works. This benchmark was developed to assess the auto-vectorization capabilities of compilers. Therefore, we only use those loop nests in the serial code related experiments. The second benchmark suite that we collect loop nests from is Polybench[21]. This suite consists of 30 benchmarks that perform numerical computations used in various domains such as linear algebra computations, image processing, physics simulation, etc. We use Polybench for experiments involving both serial and auto-parallelized code. We use the two largest datasets from Polybench to create our ML dataset. In our experience, the variance of both the hardware performance counter values and the most suited code optimizer for the loop nests across the two datasets, was enough to treat them as two different loop nests. This variance can be attributed to two main reasons. First, a different set of optimizations being performed by the optimizers based on the built-in analytical models/heuristics that drive those optimizations, since properties like loop trip counts usually vary across datasets. Second, the performance across datasets on an architecture with a memory hierarchy, where the behavior of memory may change on one or more levels. This analysis was required to prevent the ML algorithms from *overfitting*.

Code Optimizer	Version	Flags (Auto-Parallelization flags)	Auto-Parallelization
clang (LLVM)	6.0.0	-Ofast -march=native	No
gcc (GNU)	5.4.0	-Ofast -march=native	No
icc (Intel)	18.0.0	-Ofast -xHost (-parallel)	Yes
polly	6.0.0	-O3 -march=native -polly -polly-vectorizer=stripmine -polly-tiling (-polly-parallel)	Yes

Table 1: Candidate Code Optimizer and their flags

### 2.3 Experimental Platforms

For the experiments, we used two recent Intel architectures. The first architecture is a four-core Intel Kaby Lake Core i7-7700K. This architecture supports Intel’s SSE, AVX and AVX2 SIMD instruction set extensions. The second architecture we use is a two sixteen-core Intel Skylake Xeon Gold 6142. The Skylake architecture supports two more SIMD instruction set extensions, i.e., AVX-512CD and AVX-512F than the Kaby Lake architecture. For the auto-parallelization related experiments, only one thread is mapped per core.

We skip dynamic instruction count as a feature and normalize the rest of the hardware performance counters in terms of *per kilo instructions* (PKI). We exclude loop nests that have low value for crucial hardware performance counters such as instructions retired. From our experiments, we discovered two interesting correlations among hardware performance counters and the characteristic behavior of the loop nests. First, the hardware performance counters values from Kaby Lake architecture (after disabling loop transformations and vector code generation) were sufficient to get well trained ML model to make predictions for a similar architecture like the Skylake architecture. Second, for predicting the most suited candidate for serial code and for the auto-parallelized code for a loop nest, the same set of hardware performance counters, collected from profiling a serial version, can be used to train the ML model and achieve satisfactory results.

### 2.4 Machine Learning Model Evaluation

For training and evaluating our Machine Learning model, we use Orange[8]. We use Random Forest (RF) as the classifier for all the experiments. We randomly partition our dataset into Training dataset (75%) and Validation dataset (25%). The training dataset allows us to train and tune the ML models. We evaluate our trained models on Accuracy and Area Under Curve (AUC). Whereas, the validation dataset is a set of unseen loop nests that we use to make predictions. For serial code experiments, there are 209 instances (loop nests) in the training dataset and 69 instances in the validation dataset. For auto-parallelized code experiments, there are 147 instances in the training dataset and 49 instances in the validation dataset. The predicted optimizer’s execution time as compared to that of the most suited optimizer’s execution time will be same in case of correct predictions and higher in case of mispredictions.

We repeat our ML experiments thrice in order to validate our results, i.e., we randomly split the dataset, train new ML models and then make the predictions. We take into account the unique instances from the three validation datasets for measurements. Therefore, the number of instances differ between similar experiments.

## 3 Experimental Analysis

For evaluating the results, we calculate the speedup of ML predictions over candidate code optimizers, i.e., the speedup obtained if the code optimizer recommended by the ML model was used to optimize loop nests instead of a candidate.

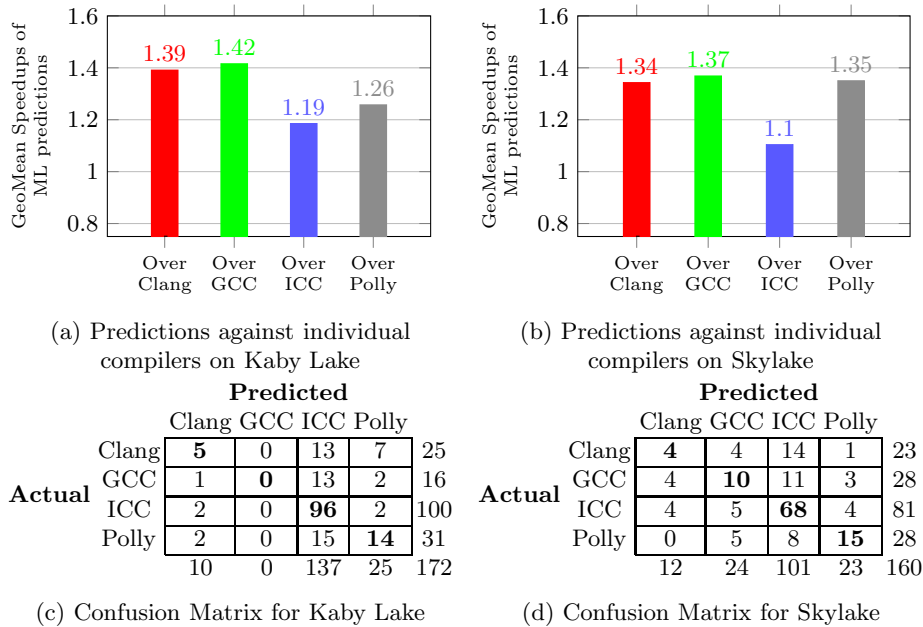


Fig. 1: Speedup of Predictions for Serial Code

### 3.1 Predicting the Most Suited Code Optimizer for Serial Code

Fig. 1a and Fig. 1b show the results for the performance gains from the predictions for the Kaby Lake and Skylake architectures, respectively. These predicted gains can be viewed as the achievable headroom for each compiler. On the validation dataset, RF classifier predicted with an overall accuracy of 67% for Kaby Lake and 61% for Skylake as shown in the confusion matrices in Fig. 1c and Fig. 1d respectively.

Across both architectures, Intel compiler performs well on majority of the loop nests. Therefore, the Majority Classifier predicted ICC with 58% overall accuracy for Kaby Lake and 50% overall accuracy for Skylake. The distribution of performance of the ML predictions compared to ICC, the maximum performance gain on a loop nest was 27x, whereas the maximum slowdown was 0.2x.

### 3.2 Predicting the Most Suited Code Optimizer for Auto-Parallelized Code

For the auto-parallelization experiments, there are only two candidates: ICC and Polly. The RF classifier predicted with an overall accuracy of 85% for Kaby Lake and 72% for Skylake as shown in Fig. 2c. Since the validation dataset was well balanced for the two targets, the Majority Classifier produced an overall accuracy of 64% for Kaby Lake and 50% for Skylake. Based on the distribution




Fig. 2: Speedup of Predictions for Auto-Parallelized Code

of performance of the ML predictions, when compared to ICC, the maximum gain on a loop nest was 91x whereas the maximum slowdown was 0.09x.

## 4 Overall Analysis and Discussion

The performance gain from the ML predictions over the candidate code optimizers range from 1.10x to 1.42x for the serial code and from 1.30x to 1.71x for the auto-parallelized code across two multi-core architectures. Counters related to Cycles Per Instruction (CPI), D-TLB, memory instructions, cache performance (L1, L2 and L3) and stall cycles were crucial indicators of the inherent behavior of the loop nests.

On analyzing the validation datasets for serial code experiments, we found that on an average for 95% of the loop nests, there was at least 5% performance difference between the most suited code optimizer and the worse suited code optimizer. For auto-parallelized code experiments, on an average for 91.5% of the loop nests, there was at least 5% performance difference between the most suited code optimizer and the worse suited code optimizer.

On the other hand, for the serial code experiments, for 68% of the loop nests, there was at least 5% performance difference between the most suited code optimizer and the second most suited code optimizer. That suggests that for the remaining 32% of the loop nests, it would be harder to make a distinction between the most suited code optimizer and the second one. Since the ML models' overall accuracy are 67% for Kaby Lake and 61% for Skylake, we can infer that they are doing very well on the loop nests that have a clear distinction about the most suited code optimizer.

---

**Acknowledgments.** This work was supported by NSF award XPS 1533926.

## References

1. R. Allen and K. Kennedy. Automatic translation of fortran programs to vector form. *ACM Trans. Program. Lang. Syst.*, 9(4):491–542, Oct. 1987.
2. A. H. Ashouri and et. al. MiCOMP: Mitigating the compiler phase-ordering problem using optimization sub-sequences and machine learning. *ACM Transactions on Architecture and Code Optimization (TACO)*, 14(3):29, 2017.
3. U. Bondhugula and et. al. Automatic transformations for communication-minimized parallelization and locality optimization in the polyhedral model. In *International Conference on Compiler Construction*, pages 132–146. Springer, 2008.
4. D. Callahan, J. Dongarra, and D. Levine. Vectorizing compilers: A test suite and results. In *Proceedings of the 1988 ACM/IEEE Conference on Supercomputing, Supercomputing '88*, pages 98–105, Los Alamitos, CA, USA, 1988. IEEE Computer Society Press.
5. R. Cammarota and et. al. Optimizing program performance via similarity, using a feature-agnostic approach. In *Advanced Parallel Processing Technologies*, pages 199–213, Berlin, Heidelberg, 2013. Springer.
6. J. Cavazos and et. al. Rapidly selecting good compiler optimizations using performance counters. In *Code Generation and Optimization, 2007. CGO'07. International Symposium on*, pages 185–197. IEEE, 2007.
7. A. Darte, Y. Robert, and F. Vivien. *Scheduling and automatic Parallelization*. Springer Science & Business Media, 2012.
8. J. Demšar and et. al. Orange: data mining toolbox in python. *The Journal of Machine Learning Research*, 14(1):2349–2353, 2013.
9. G. Fursin and et. al. Milepost GCC: Machine learning enabled self-tuning compiler. *International journal of parallel programming*, 39(3):296–327, 2011.
10. Z. Gong and et. al. An empirical study of the effect of source-level loop transformations on compiler stability. *Proc. ACM Program. Lang.*, 2(OOPSLA):126:1–126:29, Oct. 2018.
11. T. Grosser, A. Groesslinger, and C. Lengauer. Polly - Performing Polyhedral Optimizations on a Low-level Intermediate Representation. *Parallel Processing Letters*, 22(04), 2012.
12. K. Kennedy and J. R. Allen. *Optimizing Compilers for Modern Architectures: A Dependence-based Approach*. Morgan Kaufmann Publishers Inc., 2002.
13. W. Li and K. Pingali. A singular loop transformation framework based on non-singular matrices. In *International Workshop on Languages and Compilers for Parallel Computing*, pages 391–405. Springer, 1992.
14. A. W. Lim, G. I. Cheong, and M. S. Lam. An Affine Partitioning Algorithm to Maximize Parallelism and Minimize Communication. In *Proceedings of the 13th International Conference on Supercomputing, ICS '99*, pages 228–237, New York, NY, USA, 1999. ACM.
15. A. W. Lim and M. S. Lam. Maximizing Parallelism and Minimizing Synchronization with Affine Partitions. *Parallel Comput.*, 24(3-4):445–475, May 1998.
16. A. W. Lim, S.-W. Liao, and M. S. Lam. Blocking and Array Contraction Across Arbitrarily Nested Loops Using Affine Partitioning. In *Proceedings of the Eighth ACM SIGPLAN Symposium on Principles and Practices of Parallel Programming, PPOPP '01*, pages 103–112, New York, NY, USA, 2001. ACM.
17. S. Maleki and et. al. An evaluation of vectorizing compilers. In *2011 International Conference on Parallel Architectures and Compilation Techniques*, pages 372–382, Oct 2011.



18. Padua, Kuck, and Lawrie. High-speed multiprocessors and compilation techniques. *IEEE Transactions on Computers*, C-29(9):763–776, Sept 1980.
19. D. A. Padua and M. Wolfe. Advanced Compiler Optimizations for Supercomputers. *Commun. ACM*, 29(12):1184–1201, 1986.
20. Polly: LLVM Framework for High-Level Loop and Data-Locality Optimizations. <http://polly.llvm.org>.
21. PolyBench/C 4.1. <http://web.cse.ohio-state.edu/~simspouchet/software/polybench/>.
22. K. Stock, L.-N. Pouchet, and P. Sadayappan. Using machine learning to improve automatic vectorization. *ACM Transactions on Architecture and Code Optimization (TACO)*, 8(4):50, 2012.
23. G. Tournavitis and et. al. Towards a holistic approach to auto-parallelization: Integrating profile-driven parallelism detection and machine-learning based mapping. In *Proceedings of the 30th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI '09*, pages 177–187, New York, NY, USA, 2009. ACM.
24. Z. Wang and M. F. O’Boyle. Mapping parallelism to multi-cores: A machine learning based approach. In *Proceedings of the 14th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming, PPOPP '09*, pages 75–84, New York, NY, USA, 2009. ACM.
25. N. Watkinson and et. al. Using hardware counters to predict vectorization. In *Languages and Compilers for Parallel Computing (LCPC 2017)*. Springer, In Press.
26. M. J. Wolfe. *High Performance Compilers for Parallel Computing*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1995.