

A Framework for Compiler Level Statistical Analysis over Customized VLIW Architecture

Amir Hossein Ashouri, Vittorio Zaccaria, Sotirios Xydis, Gianluca Palermo and Cristina Silvano
 Dipartimento di Elettronica, Informazione e Bioingegneria, Politecnico di Milano, Italy
 (ashouri, zaccaria, xydis, gpalermo, silvano)@elet.polimi.it

Abstract—Very Long Instruction Word (VLIW) application specific processors represent an attractive solution for embedded computing, offering significant computational power with reduced hardware complexity. However, they impose higher compiler complexity since the instructions are executed in parallel based on the static compiler schedule. Therefore, finding a promising set of compiler transformations and defining their effects have a significant impact on the overall system performance. The proposed methodology provides the designer with an integrated framework to automatically (i) generate optimized application-specific VLIW architectural configurations and (ii) analyze compiler level transformations, enabling application-specific compiler tuning over customized VLIW system architectures. We based the aforementioned analysis on a Design of Experiments (DoEs) procedure that captures in a statistical manner the higher order effects among different sets of activated compiler transformations. Applying the proposed methodology onto real-case embedded application scenarios, we show that (i) only a limited set of compiler transformations exposes high confidence level (over 95%) in affecting the performance and (ii) using them we could be able to achieve gains between (16-23)% in comparison to the default optimization levels.

I. INTRODUCTION

Embedded systems design traditionally exploits the knowledge of the target domain, e.g. telecommunication, multimedia, home automation etc., to customize the HW/SW coefficients found onto the deployed computing devices. Although the functionalities of these devices are differed, the computational structure and design are tightly connected with the platform in which they rely on. Platform-based design has been proposed as a promising alternative for designing complex systems by redefining the problem of designing into that of finely tuning specific parameters of the platform template.

The scientific and commercial urge to use VLIW technology seems to be raising again after three decades of their existence [1]; VLIW processor templates are being used especially in embedded processors, designed to perform special-purpose functions, usually for real-time or hardware acceleration. Being able to use VLIW power-saving cores in CPUs seems to be using day by day. However, the trade-offs between right parallel execution and the speedup managed by compiler instead of hardware is becoming a very complex task. VLIW can achieve far higher performance, offering high degree of Instruction Level Parallelism (ILP) with low silicon and power costs. On the one hand, architecture configurability of VLIW platforms offers significant advantages regarding portability, sizing and parameter tuning provided to the designer [1], [2]. On the other hand, it introduces a lot of complexity during optimization due to multi-objective nature of the solution space and the multi-parametric structure of the design space.

Although a significant amount of research has been conducted on exploring and optimizing VLIW architectural parameters [3] and introducing specific compiler optimization for VLIW processors [4], [5], there are limited references

regarding the analysis of the impacts of conventional compiler transformations onto VLIW architectures and how these transformations are correlating with the underlying architectural configuration. Nowadays, the existence of modular and reusable compiler tool-chains [6], [7] raises the opportunity for system designers to exploit sophisticated compiler passes and customize their compiler infrastructure accordingly. Given the large decision space provided by modern compiler infrastructures, the designer has to traverse to find the best trade-off points, thus a fine-grained and automatic characterization of the effects that each compiler transformation has onto the application's behaviour, is considered of great importance. Empirical evaluation of the effects, by simply activating and deactivating compiler passes cannot be considered adequate, since a lot of inter-transformation interactions and second order effects are neglected. Due to the complexity of characterizing the solution space, there is a necessity to extend conventional exploration approaches by applying sophisticated analysis and data-mining for extracting knowledge from statistical results [8]. The problem becomes more demanding in the embedded computing domain, which requires different optimizations related to each platform configuration customized for a specific application domain. The main contribution of this paper consists of proposing a compiler/architecture methodology that provides to the designer an integrated environment to automatically (i) generate optimized application specific architectural configurations of VLIW-based platforms and (ii) analyse the effects of compiler level transformations in a statistical manner.

The proposed methodology targets the problem of compiler/architecture co-exploration in embedded computing, and it is focused on enabling application-specific compiler tuning over customized VLIW system architectures. First, a multi-objective exploration loop targeting application-specific micro-architectural customization is applied for extracting the best VLIW architecture candidates. We utilize the newly introduced Roof-Line processor architecture model [9] for characterizing the differing architectural solutions onto various resource constraints. The optimized VLIW architectural configurations are then propagated to the compiler analysis phase in which the statistical effects of the applied compiler transformations are characterized in a fine grained manner. The developed exploration framework integrates the LLVM compiler infrastructure [6] as a source to source code transformation tool together with the VEX compiler-simulator for mapping the transformed code onto custom VLIW architecture instances. We evaluated the overall methodology (customized architecture selection and statistical compiler level analysis) by using a GSM codec application as driving use case. We show that only a limited set of compiler transformations has significant effects on optimizing performance across a set of GSM specific VLIW processors. In addition to the application specific scenario, we present results regarding multiple embedded applications onto a single VLIW instance, showing that the proposed analysis can be used to extract promising compiler transformations

regarding both an intra- and cross-application manner.

The rest of the paper is organized as follows. Section II provides a brief discussion on related work and current state of the art in the field. In Section III, we introduce the basic methodology for architecture customization and statistical compiler level analysis. Section IV presents experimental evaluation of the proposed methodology on differing customized VLIW architectures and benchmark applications. Section V summarizes of the work and concludes the paper.

II. RELATED WORK

Although we have entered the era of multi-core systems, the high degree of instruction parallelism offered by VLIW architectures seems to make them an interesting alternative for a large set of commercial embedded systems [1], [10], [11]. VLIW architectures are also emerging in the modern many-core embedded accelerator devices, i.e. KALRAY MPPA256 [12], for image and signal processing applications.

Several research works have been presented targeting to the generation of Pareto optimal VLIW architectural configurations [3], [10] by exploring the space using pre-allocated compiler sequences over differing architecture instances. Towards the same direction of VLIW architectural configuration, Wong et al [13] introduced r-VEX, a reconfigurable and extensible VLIW processor. Source code is mapped using the VEX (VLIW Example) environment [14], which forms a compilation-simulation system that targets a wide class of VLIW processor architectures, and enables compiling, simulating, analyzing and evaluating C programs [2].

In current literature, there is a lot of attention on iterative compilation and predictive compiler modelling. Aganov et al [15] have introduced iterative compiler optimization for optimizing performance. In [16] the authors developed a technique to automatically build a performance model for predicting the impact of program transformations on processor architectures. Machine learning techniques have also been proposed in [17], [18], [19] to predict the potential speedup of compiler transformed programs utilizing code features provided by static program analysis. However, there is a lack of comprehensive analysis regarding the impact of applying differing conventional compiler transformations on customized VLIW architectures. Although, in VLIW compilation infrastructures [14] there are available batch compiler optimization modes, fine-grained analysis of compiler effects for VLIW architectures and its relation with architecture customization is not adequately targeted.

III. METHODOLOGY FOR COMPILER ANALYSIS OF CUSTOMIZED VLIW ARCHITECTURES

In this section, we describe the proposed methodology for compiler analysis of customized VLIW architectures. The proposed methodology consists of two phases: (i) Customized VLIW architecture selection and (ii) Statistical analysis of compiler transformations. From a high level point of view, we first generate a set of promising VLIW architectural candidates that tailors to the characteristics of the target application, optimizing on the performance-intensity trade-off curve with respect to the overall hardware allocated resource. Then, statistical analysis of distributions generated over the compiler transformation space is performed on the set of these selected customized VLIW solutions. This enables the designer to characterize the effects of each compiler transformation in both an architecture specific manner and a cross-architecture manner.

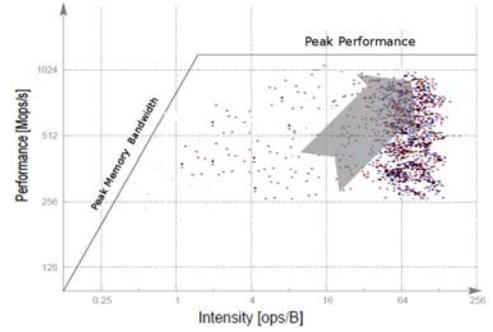


Fig. 1. Example of application of the Roof-line performance model.

We used the Roof-Line performance model [9] as the basis for both generating the custom architecture configurations and characterizing the effect of the compiler passes. Roof-Line relates processor performance to off-chip memory traffic. It [9] characterizes processor architectures in a two-dimensional space: performance (Mops/sec) vs. operational intensity (ops/Byte). *Operational intensity* is defined as operations per byte of DRAM traffic, defining total byte accessed as those bytes that go to the main memory after been filtered by the cache hierarchy. The advantage of using Roof-Line model is twofold: (i) it provides the designer with an intuitive visual metric for fast evaluation of the architectural optimality of the configuration and (ii) it is useful to characterize the impact of applied compiler transformations onto a specific architecture. For example, Figure 1 presents the Roof-Line model of a specific VLIW configuration and the superposition of application configurations derived by an experimental campaign of 4K different compiler parameter combinations. A general trend (highlighted by the arrow in Figure 1) can be easily detected towards higher performance and operational intensity points. Given this visual representation, a designer can detect promising compiler passes to be applied.

A custom exploration and analysis framework (Figure 2) has been developed based on the integration of open source tools to implement the proposed methodology. Specifically, we used Multicube Explorer [20] as the central DSE engine. Given the architectural and compiler design space descriptions,

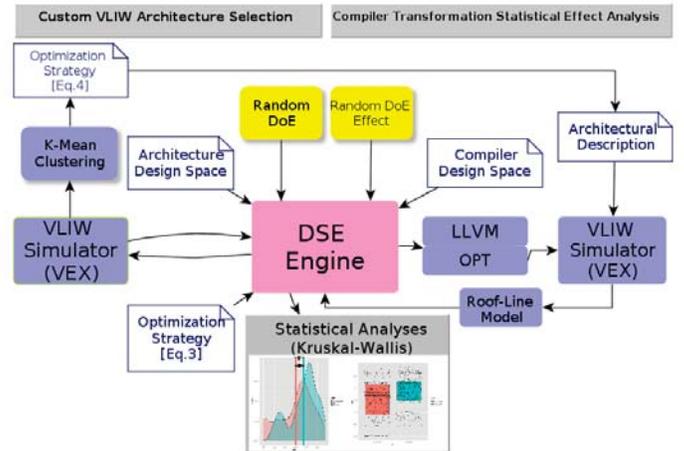


Fig. 2. Tool-chain implementing the proposed methodology

it automatically generates configuration vectors according to the specified DoE – random DoE during the phase of custom VLIW architecture selection and random effect DoE during the compiler transformation analysis phase. The LLVM compiler infrastructure [6] is integrated within the framework – LLVM C front-end and opt tool – as a source to source transformation tool to apply the compiler transformations instructed by the DSE engine. The transformed code of the application is mapped onto the VLIW processor by using the VEX [14] VLIW compiler-simulator tool, which is used for both generating different VLIW architectural configurations and mapping code onto these custom VLIW processors. Custom scripts have been developed to evaluate each examined configuration according to the Roof-Line model. Statistical analysis and visualization of results are performed by using the R statistical language [21].

A. Custom VLIW Architecture Selection

Application-specific customization of architecture’s parameters is one of the early system design optimization phases for defining platform configurations that meet the desired performance specifications. Given the large number of parameters and the delay required for simulating each possible configuration, the task of optimal micro-architectural parameter selection forms an extremely challenging exploration problem that becomes quickly intractable, regarding the time required for exhaustive evaluation. Several research works utilizing well-known meta-heuristics [3], [22] have been already proposed for generating the Pareto optimal sets of the aforementioned optimization problem.

In this paper, however, we slightly shift the focus of exploration from delivering the optimal set of architectural configurations to discover custom architecture configurations that do not correspond to the boundaries of Pareto regions, i.e. very low cost architectures with very poor performance or very expensive architectures that deliver very high gains regarding performance. Thus, in this paper we invoke a relaxed optimization search strategy that is based on a random sampling of the targeted design space rather than on a optimization oriented strategy, e.g. simulated annealing or NSGA-II genetic optimization [22] etc.

Table I shows the micro-architectural design space, Ω , considered for the custom VLIW architecture selection phase. In the first step, we randomly sample the Ω design space. Each explored solution is stored in the database of explored solutions, X after being characterized according to the performance and operational intensity metrics defined within the Roof-Line model, where:

$$Performance(\mathbf{x}) = \frac{\#Operations(\mathbf{x})}{\#NumCycles(\mathbf{x}) \times ClkFreq(\mathbf{x})} \quad (1)$$

$$Intensity(\mathbf{x}) = \frac{\#Operations(\mathbf{x})}{\#CacheMisses(\mathbf{x}) \times CacheLineSize(\mathbf{x})} \quad (2)$$

After the formation of the X , we are interested in finding those explored architectures that maximize the performance and operational intensity of the application while using minimum computational and memory resources. In order to extract the desired architectural configurations, we perform Pareto filtering on the solution space defined by X , by considering the following multi-objective optimization problem:

TABLE I. VLIW MICROARCHITECTURAL DESIGN SPACE

Parameters	Values (Integer Range)
lg2CacheSize	[11,30]
lg2Sets	[0,3]
lg2LineSize	[5,9]
lg2lCacheSize	[11,30]
lg2lCacheSets	[0,3]
lg2lCacheLines	[5,9]
ClkFreq	[300,500]
NumCaches	[1,2]
IssueWidth	[1,16]
NumAlus	[1,16]
NumMuls	[1,4]
RegisterFile	[32,128]
BranchRegister	[32,128]

$$\min_{\mathbf{x} \in \Omega} \left[\begin{array}{c} 1 \\ \overline{Performance(\mathbf{x})} \\ 1 \\ \overline{Intensity(\mathbf{x})} \\ \#CompResources(\mathbf{x}) \\ \#MemResources(\mathbf{x}) \end{array} \right] \quad (3)$$

where computational resources are (i) number of ALUs and (ii) number of multipliers, while memory resources are (i) data cache size, (ii) instruction cache size and (iii) register file size. Although, in Eq. 3 we present the unconstrained version of the target optimization problem, our exploration infrastructure permits also the inclusion of arbitrary constraints either on the objectives itself or on specific parameter combinations that the designer has a-priori evaluated as not interesting.

The outcome of the optimization procedure defined in Eq. 3 is a Pareto surface, X_p , of the explored X , thus exhibiting a large number of VLIW architectural configurations. In order to restrict the number of VLIW configuration that will be characterized as the representative customized VLIW solutions that will be propagated to the statistical compiler analysis phase, we perform a clustering on the performance - intensity solution space. We used k-means [23] clustering for the aforementioned procedure, with a configurable number of clusters, k, decided by the designer. The clustering procedure partitions the X_p solution space into k regions of interest, $X_p^{c_i}$, e.g. region of high intensity and high performance, or region of low intensity and high performance etc. Eventually, each cluster should deliver one representative VLIW architecture, that forms the optimal solution within the cluster. We define this optimal solution per cluster as the architectural configuration that minimizes area cost of the processor while maximizing both the metrics of performance and operational intensity. In order to extract this optimal configuration from each cluster, we iteratively apply the following single-objective minimization problem in every $X_p^{c_i}$ generated by the k-means clustering:

$$\min_{\mathbf{x} \in X_p^{c_i}} \frac{Area(\mathbf{x})}{Performance(\mathbf{x}) \times Intensity(\mathbf{x})} \quad (4)$$

For the calculation of the area cost in Eq. 4, the area model provided by the McPAT [24] micro-architecture framework has been used, assuming a process technology of 90 nm.

B. Compiler Transformation Statistical Effect Analysis

This second phase of the proposed methodology receives as input the custom VLIW architectures generated as described in the previous section, and for each of them it evaluates the statistical effects of the compiler transformations in a fine grained manner. In this research work we focus on 15 of the compiler passes supported by LLVM (see Table II, [6]).

TABLE II. SELECTED COMPILER TRANSFORMATIONS FROM LLVM [6] FRAMEWORK

Compiler Transformation	Abbreviation	Short Description
Constant Propagation	constprop	Instructions involving only constant operands are replaced with a constant value and propagated
Dead Code Elimination	dce	It checks instructions that were used by removed instructions to see if they are newly dead
Function Integration/Inlining	inline	Bottom-up inlining of functions into callees
Combine Redundant Instruction	instcombine	Combine instructions to form fewer, simple instructions. This pass does not modify the CFG and is where algebraic simplification happens
Loop Invariant Code Motion	licm	Attempting to remove as much code from the body of a loop as possible. It does this by either hoisting code into the pre-header block, or by sinking code to the exit blocks if it is safe
Loop Strength Reduction	loop-reduce	It performs a strength reduction on array references inside loops that have as one or more of their components the loop induction variable
Rotates Loops	loop-rotate	A simple loop rotation transformation
Unroll Loops	loop-unroll	This pass implements a simple loop unroller
Unswitch Loops	loop-unswitch	This pass transforms loops that contain branches on loop-invariant conditions to have multiple loops
Promote Memory To Register	mem2reg	It promotes memory references to be register references.
Memorycopy Optimizations	memcpyopt	It performs various transformations related to eliminating memcpy calls, or transforming sets of stores into memset's
Reassociate Expressions	reassociate	It reassociates commutative expressions in an order that is designed to promote better constant propagation
Scalar Replacement of Aggregates	scalarrpl	It breaks up alloca instructions of aggregate type (structure or array) into individual alloca instructions for each member if possible
Sparse Conditional Constant Propagation	sccp	It assumes values are constant and Basic Blocks are dead unless proven otherwise. It proves values to be constant, and replaces them with constants and Proves conditional branches to be unconditional
Simplify the Control Flow Graph	simplycfg	Performs dead code elimination and basic block merging

As a first step in our analysis, we have to determine a reasonable amount of samples to produce a robust analysis of the main effects associated with the 15 compiler parameters. In the following, each configuration of these compiler parameters (or **options set**) will be defined as a vector of 15 values, where each value represents a compiler pass option.

To accommodate our goal, we defined a randomized design of experiments $D_N(p)$ for each compiler parameter p . $D_N(p)$ is a list of *options sets*:

$$D(p) = [o_{1+}, o_{1-}, o_{2+}, o_{2-}, \dots, o_{N+}, o_{N-}] \quad (5)$$

where o_{n+} corresponds to the n -th random *option set* in which compiler pass $p \in \{OFF, ON\}$ is set to its maximum value (ON) while all the others compiler passes are randomly chosen. In a dual way, o_{n-} is equal to o_{n+} except that p assumes its minimum value (OFF).

By applying this DoE, we can easily measure how much the impact of the transition ($- \rightarrow +$) for parameter p impacts (in average over all the considered *options sets*) on the performance without requiring a full-factorial design. As an example, Figure 3 shows the performance distributions (density) generated by activating and deactivating the 'licm' and 'reassociate' compiler transformations for a GSM codec application. It can be observed that while the activation of 'licm' has a clear positive effect on performance – the median is shifted towards higher performance values. This is not the case for the 'reassociate' transformation, since the activation and deactivation distributions have almost the same shape and density, thus not permitting the designer to recognize a clear trend.

As the second step, for each options set in $D(p)$ we evaluate the vector of *performance responses* with the actual architecture synthesis after the compilation and simulation of the target application. We consider the hypothesis whether the mean of the performance given by the *options sets* where p was minimum (or *off*) is *different* from the mean where p was maximum (or *on*). In practice, this is framed as a **null-hypothesis statistical test**, which, given the *non-parametric* (or non-gaussian) nature of the underlying distributions¹, cannot be assessed with as a simple ANOVA but, instead, with a *Kruskal-Wallis* test [25]. To complete the hypothesis test, the designer sets an acceptance ratio of $p - value\%$ meaning

¹Since the distributions are built based on empirical/experimental data, the distribution is considered in general non-parametric

that the probability of 'measuring' different means when the underlying distributions are equal (or the chance of a false positive) is less than 5%.

IV. EXPERIMENTAL EVALUATION

In this section, we experimentally assess the proposed methodology. We consider the GSM codec embedded application as driving use case, automatically generating four representative application specific architectures after applying the custom VLIW architecture selection. We use these VLIW architectures for statistically analysing the effects of compiler transformations across several VLIW configurations. Furthermore we analyse the compiler transformation effects in a cross application manner, by considering a larger set of embedded applications mapped onto a default (non application specific) VLIW processor configuration.

We apply the overall proposed methodology by considering the GSM codec as the driving application. We apply the custom VLIW architecture selection phase to generate optimized representative VLIW architectures in an application specific manner. The considered architectural design space is summarized in Table I. We configure the search procedure to randomly generate and evaluate 30K configurations, by using a uniform sampling over the targeted configuration space (Table I). Applying the multi-objective optimization problem defined in Eq. 3 over the 30K solutions, the Pareto surface of the configurations that maximize performance and operational intensity while minimizing resources is generated. Without

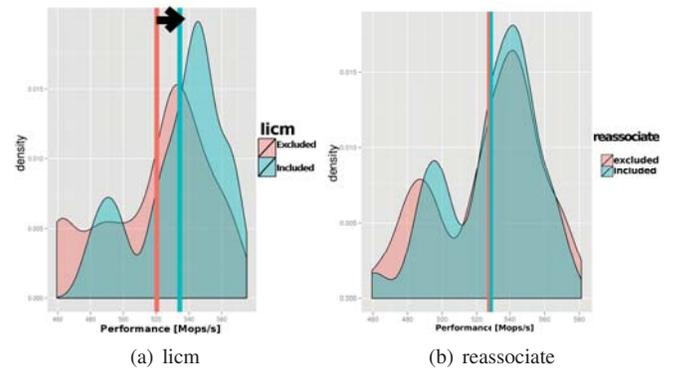


Fig. 3. Visualization of (a) licm's significant positive effect, (b) reassociate's no significant effect.

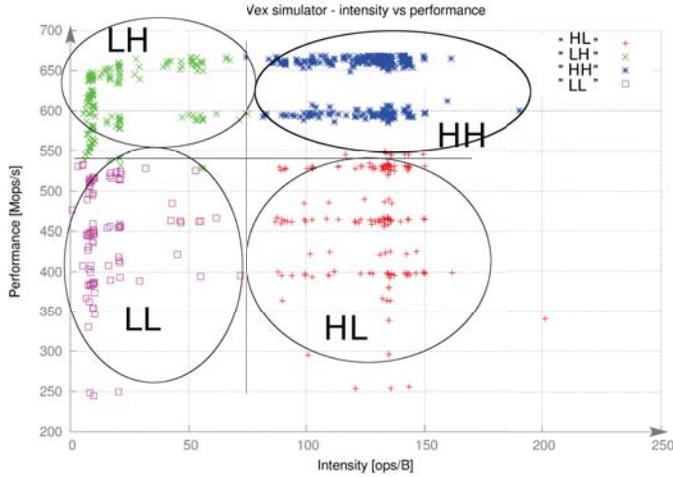


Fig. 4. Performance-Intensity exploration results: The four clusters of architectural Pareto-sets

loss of generality, we consider the generation of $k=4$ clusters over the generated Pareto surface, aiming at the generation of four GSM-specific VLIW architectures. Figure 4 shows the results of clustering of the extracted Pareto surface and its mapping onto the two-dimensional performance vs. intensity space. Each cluster has been characterized according to its position on the performance vs. intensity space as: (i) HH for the cluster placed to the high intensity and high performance region, (ii) LH for the low intensity and high performance region, (iii) LL for the low intensity and low performance region and (iv) HL for the high intensity and low performance region, respectively.

The final $k=4$ representative VLIW architectures are derived after applying within each cluster the optimization operator of Eq. 4. Table III reports the architectural configurations for each of the $k=4$ application specific VLIW architectures.

For each of the $k=4$ application specific VLIW architectures, we explore the compiler level design space, defined in Table II. We generate the non-parametric distribution of the performance and intensity for each compiler transformation by considering 500 samples per transformation. As described in Section III-B, the non-parametric distributions are analysed based on Kruskal-Wallis test to specify the statistical effects, whether or not the inclusion of a specific transformation impacts on a specific and robust manner the two considered metrics. Table IV summarizes the results of Kruskal-Wallis statistical tests for each compiler transformation over the four examined architecture configurations. As shown, four compiler passes (*inline*, *licm*, *loop-reduce* and *loop-rotate*), out of the fifteen initially considered, have a clear and significant impact on performance when activated. In addition, Figure 5, shows the confidence level for each one of the considered compiler transformations. It is shown that the four mentioned compiler transformations exhibit a high confidence level $>99\%$. Therefore, it could be implied that activating these specific transformations, the designer can be 99% confident that the effect on performance will be the same as the one predicted by the exploration.

In the second set of experiments, we perform statistical analysis in a cross-application manner. For this experimental campaign, we assume a larger set of applications (namely GSM, AES encryption engine, ADPCM codec, JPEG decoder and Blowfish block cipher). The performance of aforemen-

TABLE III. VLIW ARCHITECTURE CONFIGURATIONS

Parameters	Arch-HL	Arch-LH	Arch-HH	Arch-LL	Arch-User
lg2CacheSize	15	12	13	12	16
lg2Sets	1	3	0	1	2
lg2LineSize	7	5	5	5	5
lg2lCacheSize	16	14	16	14	16
lg2lCacheSets	1	3	3	2	2
lg2lCacheLines	6	8	7	5	6
ClkFreq	400	450	450	300	500
NumCaches	2	1	1	1	1
IssueWidth	6	6	14	9	8
NumAlus	4	6	7	3	8
NumMuls	1	4	4	14	2
MemLoad	4	3	6	5	4
MemStore	2	8	4	6	4
RegisterFile	104	100	32	76	64
BranchRegister	76	84	88	48	64

TABLE IV. SUMMARY OF KRUSKAL-WALLIS ANALYSIS ON PERFORMANCE FOR GSM-SPECIFIC VLIW ARCHITECTURES

CompilerTransformation	Arch-HL	Arch-LH	Arch-HH	Arch-LL
Constprop	-	-	-	-
Dce	-	-	-	-
Inline	✓	✓	✓	✓
Instcombine	-	-	-	-
Licm	✓	✓	✓	✓
Loop reduce	✓	✓	✓	✓
Loop rotate	✓	✓	✓	✓
Loop unroll	-	-	-	-
Loop unswitch	-	-	-	-
Mem2reg	-	-	-	-
Memcpyopt	-	-	-	-
Reassociate	-	✓	✓	✓
Scalarrepl	-	-	-	-
Sccp	-	-	-	-
Simplifycfg	-	-	-	-

tioned applications has been evaluated by considering a user specified VLIW architecture, Arch-User, defined in the last column of Table III. For each benchmark the compiler transformation statistical effect analysis (section III-B) is applied, considering distributions of 500 samples per compiler transformation. Table V summarizes in an aggregated manner the results of the Kruskal-Wallis analysis by considering in each case a confidence level $\geq 5\%$. For the specific setup, we observe that there is a set of four compiler parameters (*licm*, *loop reduce*, *loop rotate* and *mem2reg*), with significant effects on performance and with a high confidence level over all the examined application use cases. Furthermore, examining each application in isolation, the designer can derive which are the compiler parameters that need to be pre-allocated, thus reducing significantly the design-time required to optimize the performance of the targeted application during iterative compilation steps. In Figure 6, we report the normalized

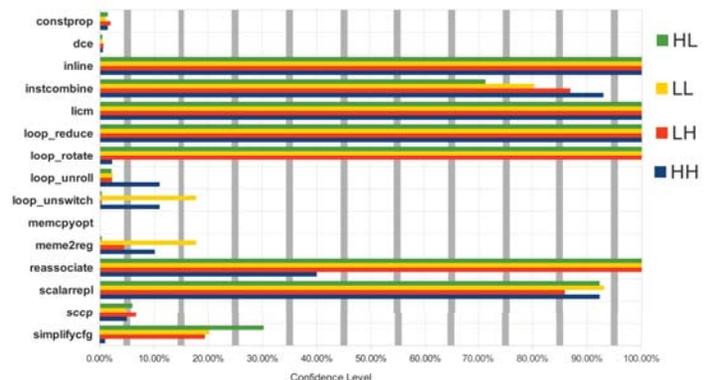


Fig. 5. Confidence level characterization of compiler transformations regarding the effect on performance for each one of the GSM specific VLIW architectures, resulted after Kruskal-Wallis statistical test.

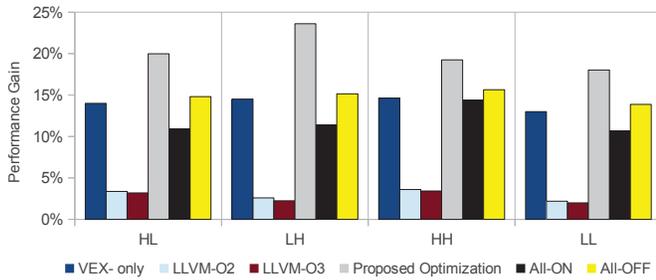


Fig. 6. Performance gain comparing to LLVM-O1 optimization level in GSM benchmark.

TABLE V. KRUSKAL-WALLIS ANALYSIS ON PERFORMANCE FOR MULTIPLE APPLICATIONS

CompilerTransformation	GSM	AES	ADPCM	JPEG	Blowfish
Constprop	-	-	-	-	-
Dce	-	-	-	-	-
Inline	✓	-	✓	✓	-
Instcombine	✓	-	✓	✓	✓
Licm	✓	✓	✓	✓	✓
Loop reduce	✓	✓	✓	✓	✓
Loop rotate	✓	✓	✓	✓	-
Loop unroll	-	-	-	-	-
Loop unswitch	-	-	-	-	-
Mem2reg	✓	✓	✓	✓	✓
Memcpyopt	-	-	-	-	-
Reassociate	✓	-	-	-	-
Scalarrepl	✓	-	-	-	✓
Sccp	-	-	-	-	-
Simplifycfg	-	-	-	-	-

speedup gains achieved by activating the compiler transformations proposed by our methodology in comparison with several well-known compilation strategies. It is shown that the proposed methodology delivered speedup gains between 16-23% in all the examined cases.

V. CONCLUSIONS AND FUTURE WORK

This paper presents a new customization and analysis methodology for compiler/architecture co-exploration of VLIW-based platform design. The proposed methodology provides the designer with an integrated framework to automatically (i) generate optimized application specific VLIW architectural configurations and (ii) analyse in a fine-grained manner the effects of compiler level transformations regarding the performance and operational intensity trade-offs. Being focused more on the analysis, we showed 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. Future work is aligned with our strong belief that the proposed analysis methodology can be further exploited for automated performance optimization through iterative compilation and architecture specialization.

REFERENCES

- [1] J. Fisher, P. Faraboschi, and C. Young, "Vliw processors: Once blue sky, now commonplace," *Solid-State Circuits Magazine, IEEE*, vol. 1, no. 2, pp. 10–17, 2009.
- [2] J. A. Fisher, P. Faraboschi, and C. Young, *Embedded computing: a VLIW approach to architecture, compilers and tools*. Morgan Kaufmann, 2004.
- [3] G. Ascia, V. Catania, M. Palesi, and D. Patti, "A system-level framework for evaluating area/performance/power trade-offs of vliw-based embedded systems," in *Design Automation Conference, 2005. Proceedings of the ASP-DAC 2005*, vol. 2, Jan., pp. 940–943 Vol. 2.
- [4] J. A. Fisher, "Trace scheduling: A technique for global microcode compaction," *IEEE Trans. Computers*, vol. 30, no. 7, pp. 478–490, 1981.
- [5] W.-M. W. Hwu, S. A. Mahlke, W. Y. Chen, P. P. Chang, N. J. Warter, R. A. Bringmann, R. G. Ouellette, R. E. Hank, T. Kiyohara, G. E. Haab *et al.*, "The superblock: an effective technique for vliw and superscalar compilation," in *Instruction-Level Parallelism*. Springer, 1993, pp. 229–248.
- [6] (2013) The LLVM website. [Online]. Available: <http://www.llvm.org/>
- [7] Rose compiler infrastructure. [online], available: [Online]. Available: <http://rosecompiler.org/>
- [8] D. Fenacci, B. Franke, and J. Thomson, "Workload characterization supporting the development of domain-specific compiler optimizations using decision trees for data mining," in *Proceedings of the 13th International Workshop on Software & Compilers for Embedded Systems*, ser. SCOPE '10. New York, NY, USA: ACM, 2010, pp. 5:1–5:10.
- [9] S. Williams, A. Waterman, and D. Patterson, "Roofline: an insightful visual performance model for multicore architectures," *Communications of the ACM*, vol. 52, no. 4, pp. 65–76, 2009.
- [10] D. Saptono, V. Brost, F. Yang, and E. Prasetyo, "Design space exploration for a custom vliw architecture: Direct photo printer hardware setting using vex compiler," in *Proceedings of the 2008 IEEE International Conference on Signal Image Technology and Internet Based Systems*, ser. SITIS '08. Washington, DC, USA: IEEE Computer Society, 2008, pp. 416–421.
- [11] P. Faraboschi and F. Homewood, "St200: A vliw architecture for media-oriented applications," in *Microprocessor Forum 2000. San Jose, CA, 2000*.
- [12] (2012) Core architecture. Kalray. Orsay, France. [Online]. Available: <http://www.kalray.eu/technology/>
- [13] S. Wong, T. Van As, and G. Brown, " ρ -vex: A reconfigurable and extensible softcore vliw processor," in *ICECE Technology, 2008. FPT 2008. International Conference on*. IEEE, 2008, pp. 369–372.
- [14] Hewlett-packard laboratories. vex toolchain. [online], available: [Online]. Available: <http://www.hpl.hp.com/downloads/vex/>
- [15] 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.
- [16] J. Cavazos, C. Dubach, F. Agakov, E. Bonilla, M. F. O'Boyle, G. Fursin, and O. Temam, "Automatic performance model construction for the fast software exploration of new hardware designs," in *Proceedings of the 2006 international conference on Compilers, architecture and synthesis for embedded systems*. ACM, 2006, pp. 24–34.
- [17] C. Dubach, J. Cavazos, B. Franke, G. Fursin, M. F. O'Boyle, and O. Temam, "Fast compiler optimisation evaluation using code-feature based performance prediction," in *Proceedings of the 4th international conference on Computing frontiers*. ACM, 2007, pp. 131–142.
- [18] C. Dubach, J. Cavazos, B. Franke, G. Fursin, M. F. P. O'Boyle, and O. Temam, "Fast compiler optimisation evaluation using code-feature based performance prediction," in *Conf. Computing Frontiers, 2007*, pp. 131–142.
- [19] H. Leather, E. V. Bonilla, and M. F. P. O'Boyle, "Automatic feature generation for machine learning based optimizing compilation," in *CGO, 2009*, pp. 81–91.
- [20] Multicube explorer. [Online]. Available: <http://m3explorer.sourceforge.net/>
- [21] R. I. R Gentleman. (2012) R statistical tool. University of Auckland. [Online]. Available: <http://www.r-project.org/>
- [22] G. Palermo, C. Silvano, S. Valsecchi, and V. Zaccaria, "A system-level methodology for fast multi-objective design space exploration," in *Proceedings of the 13th ACM Great Lakes symposium on VLSI*. ACM, 2003, pp. 92–95.
- [23] T. Kanungo, D. Mount, N. Netanyahu, C. Piatko, R. Silverman, and A. Wu, "An efficient k-means clustering algorithm: analysis and implementation," *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, vol. 24, no. 7, pp. 881–892, 2002.
- [24] S. Li, J. H. Ahn, R. D. Strong, J. B. Brockman, D. M. Tullsen, and N. P. Jouppi, "Mcpat: an integrated power, area, and timing modeling framework for multicore and manycore architectures," in *Microarchitecture, 2009. MICRO-42. 42nd Annual IEEE/ACM International Symposium on*. IEEE, 2009, pp. 469–480.
- [25] N. Breslow, "A generalized kruskal-wallis test for comparing k samples subject to unequal patterns of censorship," *Biometrika*, vol. 57, no. 3, pp. 579–594, 1970.