WCCV: Improving the Vectorization
of IF-statements with Warp-Coherent Conditions

Huihui Sun
University of Münster
Münster, Germany
huihuisun@uni-muenster.de

Jie Zhao
State key Laboratory of Mathematical Engineering and
Advanced Computing
Zhengzhou, China
jie.zhao@inria.fr

Florian Fey
University of Münster
Münster, Germany
feyf@uni-muenster.de

Sergei Gorlatch
University of Münster
Münster, Germany
gorlatch@uni-muenster.de

ABSTRACT

When vectorizing programs for modern processors with SIMD extensions, IF-statements pose a challenge: existing vectorization approaches often introduce redundant computations or they resort to inefficient masked instructions.

In this paper, we introduce a new notion of *warp-coherence* for conditions that exhibit coherent run-time behavior on different lanes of a vector register. We demonstrate that warp-coherent conditions appear frequently in practice. We present Warp-Coherent Condition Vectorization (WCCV) - an approach to detecting and optimizing IF-statements with warp-coherent conditions - to efficiently vectorize programs with IF-statements while avoiding the overhead of existing methods. WCCV detects warp-coherent conditions via the affine analysis of conditional boolean expressions and branch predication of IF-statements; the runtime code generated by WCCV avoids redundant computations and masked instructions. We employ auto-tuning to find the optimal benefit-overhead ratio for WCCV. We implement WCCV on top of Region Vectorizer (RV) - an LLVM-based vectorizing compiler, and we conduct experiments on the Rodinia benchmark suite, achieving a mean speedup of 1.14× over the original vectorized and optimized code, and speedup between 0.98× and 7.02× over the scalar code on Skylake with AVX512.

CCS CONCEPTS

• **Computer systems organization** → **Single instruction, multiple data**; • **Software and its engineering** → **Compilers**; • **Computing methodologies** → *Parallel programming languages.*

KEYWORDS

Vectorization, IF-statements, SPMD-on-SIMD, Compiler Optimization, Warp-Coherence

1 INTRODUCTION

Most modern processors employ *SIMD extensions* (Single Instruction Multiple Data) that achieve high performance by simultaneously operating on the elements of vector registers. Historically, auto-vectorization approaches rely on two main techniques - loop vectorization [27] and SLP (Superword Level Parallelism) vectorization [13] - which are implemented in widely used compilers, like GCC, LLVM, Intel ICC. Unfortunately, the performance of auto-vectorization is sometimes much worse than expected in real-world applications when compiled using typical optimization flags [31]. On the other hand, using intrinsics, i.e., low-level interfaces for SIMD instructions, requires from the application developer a very deep knowledge of the target hardware architecture. The *SPMD-on-SIMD* programming model – originated by Intel in their ISPC compiler [19] and increasingly popular recently [8, 11] – allows the programmer to specify parallelism explicitly in the SPMD (Single Program Multiple Data) style, while the compiler takes care of SIMD code generation.

We aim to improve the state of the art in vectorizing programs for efficient execution on CPUs with SIMD extensions within the SPMD-on-SIMD model. In this model, an SPMD program describes parallelism in terms of kernels: multiple kernel instances (*threads*) operating on different data can be executed simultaneously. We focus on the challenge of utilizing SIMD extensions by mapping threads onto vector registers, with each SIMD lane computing one thread, while the well-understood aspect of using available cores through multi-threading is out of scope of the paper.

Inspired by the GPU terminology, we denote threads that are simultaneously executed in lock-step on different lanes of a vector register as a *warp*. The number of threads in a CPU warp is

Huihui Sun, Florian Fey, Jie Zhao, and Sergei Gorlatch

determined by both the target vector instruction set and the data type. For example, on Intel CPUs with AVX512 extensions, each extension has 32 vector registers, each of which can process 16 floating-point values with a single SIMD instruction; therefore, in this case, a warp consists of 16 threads. Whenever the execution of a warp encounters a conditional statement, there is a possibility of thread divergence that hampers performance. A typical solution to address this divergence in vectorization is that the control flow is converted into data flow using software predication techniques, such as masked instructions: both branches of an IF-statement are executed one after the other, masking out the inactive threads. Such methods are referred to as IF-conversion or linearization [1, 10, 25]. They introduce redundant computations and resort to masked instructions that are comparatively slow. Loop unswitching [29] and partial linearization [17] alleviate this problem to some extent by analyzing and retaining the IF-statements when all threads execute the same branch.

Our approach is based on the new notion of *warp-coherent conditions* that exhibit similar runtime behavior on different lanes of a vector register. We present *Warp-Coherent Condition Vectorization* (WCCV) that exploits warp-coherence in conditions and achieves optimization opportunities for programs with such IF-statements that are usually missed by existing vectorization approaches. WCCV detects warp-coherence via the affine analysis of conditional boolean expressions and branch predication of IF-statements. For warp-coherent conditions, the generated runtime code avoids the overhead of existing methods. We employ auto-tuning to find the optimal benefit-overhead ratio when applying WCCV.

Warp-coherent conditions are widely used in real-world applications. For example, Partial Differential Equation (PDE) solvers process elements guarded with boundary check conditions: binary search algorithms repeatedly evaluate a condition until the first matching element is found. The boolean expression values of these conditions change only once across the range of thread identifiers. We call such conditions *boolean-step conditions*. Computational Fluid Dynamics (CFD) simulations process elements depending on their specific types, which only differ for a few boundary elements. The boolean expression values of these conditions have higher probabilities for one branch; we call them *high-probability conditions*. In these cases, a change in branching behavior occurs rarely, resulting in taking the same branch for a large fraction of adjacent elements in such applications. As a result, both boolean-step conditions and high-probability conditions are warp-coherent conditions that provide a potential for optimization, as long as it is possible to reliably detect them.

Fig. 1 shows how these different kinds of conditions are related to each other.

WCCV uses two different methods to detect warp-coherent conditions. The first method detects boolean-step conditions based on static affine analysis. Affine analysis is usually used for analyzing memory access patterns [15, 23], while we use affine analysis to analyze the variables and expressions used in conditions in order to detect a boolean-step condition. If the static affine analysis fails, we use the second method based on the branch probability estimation to identify high-probability conditions. We develop a cost model based on the estimated branch probability and branch cost: if a
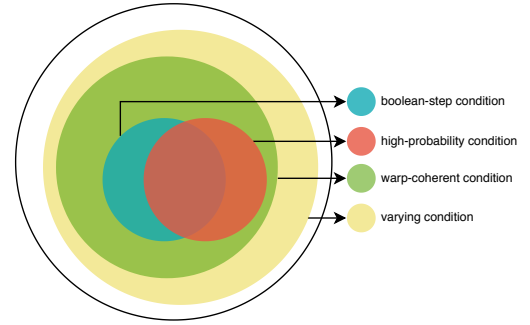

Figure 1: Relation between different kinds of conditions

certain branch is more likely to be executed and the corresponding branch cost is greater than a threshold, we treat the corresponding condition as warp-coherent. We use auto-tuning to determine the optimal thresholds for various target platforms and applications.

For each detected warp-coherent condition, WCCV inserts a runtime check of whether all threads within a particular warp execute the same branch. If the runtime check succeeds, we substitute our implementation variant for the original code, leaving out any masked instructions. If the runtime check disproves the warp-coherent condition, its threads are prone to diverging control flow, in which case we retain the standard linearization process. The main idea of WCCV is to reduce redundant computations, including masked instructions, introduced by the control flow to data flow conversion, as used in modern compilers.

Specifically, this paper makes the following contributions.

- We demonstrate that warp-coherent conditions are quite common in real-world applications, and we show that current compilers fail to vectorize such programs efficiently (Section 2).
- We propose two methods to detect warp-coherent conditions: one method detects boolean-step conditions based on static affine analysis, and the other method detects high-probability conditions based on both static branch probability estimation and auto-tuning (Section 3).
- We develop a transformation pass for detected warp-coherent conditions, including the insertion of a runtime check and the generation of the corresponding code variant before the partial linearization pass. This transformation avoids redundant computations by reducing the number of masked instructions, thereby improving the performance of the original vectorized code (Section 4).
- We implement WCCV (Section 5) at the LLVM-IR optimization level on top of Region Vectorizer (RV) [16] which implements partial linearization [17]. Our experiments on a variety of benchmarks show a mean speedup factor of 1.14 as compared to RV, and factors between 0.98 and 7.02 as compared to scalar code on Skylake with AVX512 (Section 6).

2 BASIC IDEA

In this section, we first introduce some typical example programs with warp-coherent conditions. Then, we present a motivating

code example with divergent control flow, and examine how the approaches used in modern compilers vectorize this specific example. We conclude this section by explaining how we can exploit warp-coherent conditions to vectorize this example efficiently and decrease the number of masked instructions.

2.1 Programs with Warp-coherent Conditions

Table 1 lists six programs, of which five are taken from the Rodinia benchmark suite [6] and one from AOBench [3]. The table also shows the domain of the programs as well as the so-called Berkeley Dwarf taxonomy [4]. Using WCCV, we can verify that they all have warp-coherent conditions: (1) *kmeans* and *particlefilter* have boolean-step conditions; (2) *aobench* and *cfd* have high-probability conditions; and (3) *b+tree* and *nw* have both kinds of conditions. Here we show these programs only for concept presentation, more programs with warp-coherent conditions can be found in Section 6.

Table 1: Typical programs with warp-coherent conditions

Program	Berkeley dwarf	Domain
aobench	spectral method	ray tracer
b+tree	graph traversal	graph algorithm
cfd	unstructured grid	fluid dynamics
kmeans	dense linear algebra	data mining
nw	dynamic programming	bioinformatics
particlefilter	structured grid	medical imaging

2.2 Motivating Example

Listing 1 shows a simplified kernel with an IF-statement, extracted from the *cfd* program in Table 1. Note that its control flow stays unchanged in different variants of the Rodinia benchmark suite including CUDA, OpenCL and PACXX implementations.

```
1  void compute_flux(int* elements)
2  {
3      const int tid = get_global_id(0);
4      nb = elements[tid];
5      if (nb >= 0)          /* C1 */
6      {
7          Branch 1(B1);
8      } else if (nb == -1)  /* C2 */
9      {
10         Branch 2 (B2);
11     } else if (nb == -2)  /* C3 */
12     {
13         Branch 3 (B3);
14     }
15 }
```

Listing 1: Simplified code segment from *compute_flux* in *cfd*

The Control Flow Graph (CFG) of this kernel is shown in Fig. 2a. To evaluate for each branch the probability that the runtime control flow falls in this branch, we conduct experiments by measuring the trigger ratio of each branch. As a result, branch B1 was triggered with the highest probability of 97.9%, followed by the branch B2 with 2% probability and branch B3 with the probability of only 0.1%.

GCC and LLVM. The usual approach to vectorizing this code in modern compilers such as LLVM [14] and GCC [7], is by means of *linearization* (IF-conversion) that executes all branches with masked instructions. Masked instructions prevent inactive threads (where
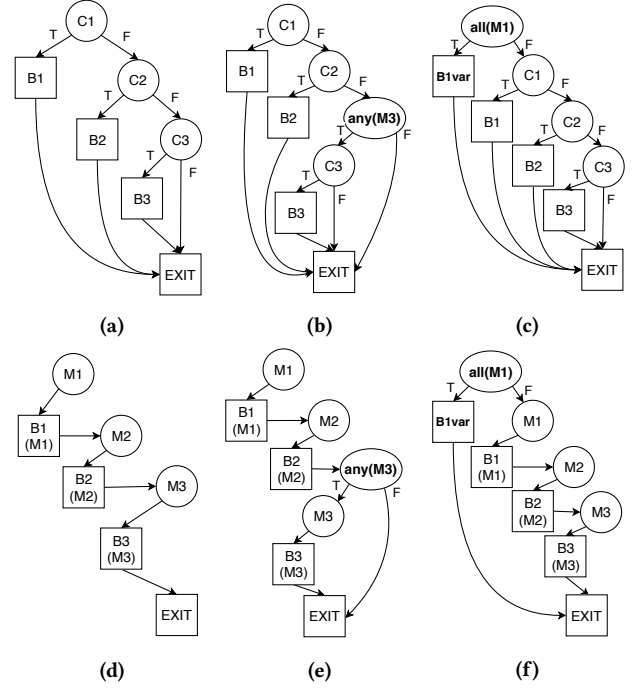

Figure 2: (a) Control flow graph of code in Listing 1, (b) With BOSCC gadget to skip B3, (c) With our runtime check and code variants of B1, (d) (e) (f) show the CFG after (partial) linearization of (a) (b) (c) respectively. M1, M2 and M3 are masks generated corresponding to C1, C2 and C3 respectively.

the condition is evaluated to false) from accessing memory and executing the corresponding branch. The control flow graph after linearization of Fig. 2a is depicted in Fig. 2d, where all branches are executed one after another with corresponding masked instructions. However, this solution is inefficient due to the significant performance penalty caused by the masked instructions. We aim to reduce the frequency of executing masked instructions using branch probability as we illustrate later in this section.

WFV and RV. The Whole Function Vectorizer [10] and Region Vectorizer [16] integrate *partial linearization* which only linearizes varying branches (different values for different threads) while preserving uniform branches (same value for different threads). In addition, RV adds a so-called BOSCC gadget that can skip the execution of a masked branch if its mask is all-false before partial linearization.

All branches of our motivating example in Listing 1 are varying. RV adds a BOSCC gadget to skip B3 since the code falls in C3 with a low probability of 0.1%. The generated CFG with BOSCC gadget is shown in Fig. 2b: the ellipse with bold *any(M3)* is the inserted BOSCC gadget that is used to check whether *M3* has at least one true bit. It retains the masked B3 branch if its predicate is true, or it skips B3 when *M3* is evaluated to all false bits. This BOSCC gadget branch is uniform, so it is preserved in the partial linearization process. The CFG after partial linearization is shown in Fig. 2e, with B1 and B2 executed one after the other with corresponding

masked instructions. If *M3* is all-false, B3 is skipped; otherwise B3 is executed with the corresponding masked instructions.

2.3 Optimization Opportunity

Masked instructions are used in the linearization phase of modern compilers such as LLVM and GCC. However, according to the Intel AVX512 manual [9], the masked instructions are 4 times slower than the ordinary instructions and their usage requires the redundant execution of branches. Therefore, vectorization using masked instructions sometimes causes performance degradation. Previous approaches incur masking overhead, which is unnecessary when the control flow is actually uniform or warp-coherent, implying that all register lanes that are used as a mask hold the same value. Our WCCV approach executes as few masked instructions as possible in order to boost the vectorization performance. We detect warp-coherent conditions and transform them to avoid masked instructions for warps whose mask is all-true (i.e., true for every thread within the warp).

For the example in Listing 1, the CFG created using WCCV is depicted in Fig. 2c, with the runtime check all(M1) and *B1var* in bold font. WCCV identifies C1 as a warp-coherent condition and it inserts a runtime check before C1. This runtime check *all(M1)* verifies whether M1 is all-true (true for every thread within the warp). If so then it branches to *B1var*; otherwise, it branches to the original condition check C1. *B1var* is a variant of B1 with uniform condition, which will not be linearized with masked instructions afterwards. The CFG after partial linearization is shown in Fig. 2f. If *M1* is all true, then *B1var* without masked instructions is executed; otherwise all three branches are executed one after another with corresponding masked instructions. WCCV differs from the BOSCC approach: the latter tries to skip the least often executed branch with masked instructions, while the former attempts to execute only the most often executed branch without masked instructions. We elaborate WCCV in the following sections.

3 DETECTING WARP-COHERENT CONDITIONS

We present two methods to help the compiler automatically detect warp-coherent conditions: 1) detection of boolean-step conditions based on affine analysis, and 2) detection of high-probability condition based on the auto-tuned branch probability estimation. The first method analyzes the variables used in the conditions using affine analysis that decides whether the condition is a boolean-step condition. If it fails, we leverage the existing branch probability pass of the LLVM compiler and use the Auto-Tuning Framework (ATF) [21] to find the optimal threshold for branch probability.

3.1 Boolean-step Conditions: Affine Analysis

We call a condition a *boolean-step condition* iff the value of the condition only changes once across the range of thread identifiers. Note that the value of the condition is considered uniform and would thus not be linearized as described in Section 2.2 if it stays unchanged in the whole range. We detect boolean-step conditions by checking that the condition's expression is affine, and if so, analyzing it to identify stepping transitions. Boolean-step conditions

consist of affine expressions and relational operators. An expression $E(i)$ of the variable $i \in \mathbb{Z}$ is affine iff it can be expressed in the form $E(i) = ai + b$ with coefficients $a, b \in \mathbb{R}$. Note that $E(i)$ is also uniform when a is equal to 0.

We extend the traditional affine analysis applied for memory access patterns [15, 23] to checking boolean-step conditions by allowing the coefficients and offsets of an affine expression to be real numbers. In our approach, we allow arbitrary expressions in a condition, including the presence of real numbers by taking into consideration various data types including float, double, integer, etc., except the integer restriction on the loop variable i. Thread identifiers are obviously affine expressions, as well as constants. The affinity of complex expressions is recursively determined by the properties of operators and their respective operands.

Fig. 3 shows how our analysis of affine expressions is used to detect boolean-step conditions. By recursively propagating the information about operands, we identify affine expressions and analyze them to detect boolean-step conditions. We denote affine expressions with subscript a, uniform expressions with subscript u, and boolean-step conditions with subscript s. Variable names are denoted with v and w. Applying unary operators like + and − to affine operands always yields affine expressions. In particular, one may interpret such operators as the sign bit of floating-point arithmetic, and the + operator would be useful when differentiating positive zero (+0) and negative zero (-0). Unary operators like (*bool*) and ! can be interpreted as logical conversion of boolean expressions, yielding step functions as results. Although boolean conversion (*bool*) is not an explicit operator, the underlying operation is important for our analysis. Additive operators applied to affine operands (including uniform operands) are affine expressions. Multiplicative operators applied to an affine operand and a constant operand (but not two general affine operands) are affine expressions. Relational operators with two affine operands yield boolean-step conditions. It may be impractical to propagate the affine state for some operators and we, therefore, view the expression as non-affine. Additionally, we also exclude conditions with variables and/or non-affine subexpressions.

Unary operators:
$$+v_a \rightarrow a \qquad\qquad -v_a \rightarrow a$$
$$(bool)v_a \rightarrow s \qquad\quad !v_a \rightarrow s$$
Additive operators:
$$v_a + w_a \rightarrow a \qquad\quad v_a - w_a \rightarrow a$$
$$v_a + w_u \rightarrow a \qquad\quad v_a - w_u \rightarrow a$$
Multiplicative operators:
$$v_a * w_u \rightarrow a \qquad\quad v_a / w_u \rightarrow a$$
$$v_a \% w_u \rightarrow a$$
Relational operators:
$$v_a == w_a \rightarrow s \qquad\quad v_a \neq w_a \rightarrow s$$
$$v_a \leq w_a \rightarrow s \qquad\quad v_a \geq w_a \rightarrow s$$
$$v_a < w_a \rightarrow s \qquad\quad v_a > w_a \rightarrow s$$

Figure 3: Detecting boolean-step conditions by affine analysis

We also make use of short-circuit evaluation of logical operators && and ‖. For example, we regard the IF-condition $if(a\&\&b)$...

as $if(a)\{if(b)...\}$. We analyze and transform these two conditions separately.

3.2 High-probability Conditions: Auto-tuned Branch Probability Estimation

In this section, we discuss detecting high-probability conditions using the statically computed probability estimation of each branch. We set an empirical threshold for which we consider high probabilities, and we auto-tune the threshold value to achieve the best performance for a particular target architecture.

We leverage the existing LLVM branch probability pass - a mid-level IR pass that statically predicts the relative probabilities of each branch in the control flow of a kernel. Besides the branch probability, the cost of the branch is also important as the overhead of the inserted checks should be amortized by the reduced masked instructions in the branch. We estimate the cost of each branch by accumulating weighted cost of instructions in the target branch. We use the auto-tuning framework ATF [21] to obtain good threshold values for the two JIT-compilation parameters – branch probability and branch cost – for different applications and target architectures.

Listing 2 demonstrates the usage of ATF for auto-tuning these two parameters for the example application *cfd* from Listing 1. When using ATF, the user annotates the source code (in this example, *cfd*'s run script – line 13 of Listing 2) as follows. First, we specify the application-specific search space by declaring in ATF the parameters *b_prob* and *b_cost* (line 1-5). We then automatically explore the search space for 10 minutes (line 9), to find well-performing parameter values using an automatic search technique. Command *run_cfd.sh* in line 7 executes the *cfd*'s run script.

```
1  #atf::tp name  "b_prob"
2           range interval<double>( 0.0, 1.0 )
3
4  #atf::tp name  "b_cost"
5           range interval<int>( 1, 500 )
6
7  #atf::run_script "./run_cfd.sh"
8
9  #atf::abort_condition duration<minutes>(10)
10
11 // cfd's run script
12 ProbMin=b_prob CostMin=b_cost ./euler3d data/fvcorr.domn.097K
```

Listing 2: The use of ATF for auto-tuning *cfd*'s parameters

4 TRANSFORMATION OF A WARP-COHERENT CONDITION

Our vectorization approach makes use of the previously explained warp-coherent conditions that ensure a high probability of converging control flow in each single warp. To fully exploit the converging control flow of warp-coherent conditions, it is necessary to introduce an additional check at runtime: we check whether the condition's corresponding mask is all-true for elements in the current vector register. If so, all threads within this warp will follow the same control flow. The runtime check is performed once per warp: due to the high probability of converging control flow, it is likely to succeed for a majority of warps. Whenever the runtime check succeeds, it is possible to use an optimized code variant without masked instructions. However, if the runtime check fails to verify converging control flow, it is necessary to use the standard code

with linearized control flow. Fortunately, the latter case rarely happens in practice because of the warp-coherence of the underlying condition.

In this section, we describe the transformation of warp-coherent conditions, including the insertion of the runtime check and the generation of the corresponding code variant. Our transformation of the example shown in Listing 1 would yield the CFG shown in Fig. 2c that differs from the original one in Fig. 2a by introducing runtime checks and code variants. The detailed explanation of the figures is presented in Section 2.

Algorithm 1: Transform Warp-Coherent Condition

Input: Conditional : BranchInst , TreatedBranchId : int
Output: transformed conditional statement

1 FirstBB := Conditional.getSuccessor(TreatedBranchId);
2 ParentBB := Conditional.getParent();
3 **if** *ParentBB dominates FirstBB* **then**
      /* Clone the treated branch: */
4   DescendantBBs := FirstBB.getDescendants();
5   CopiedFirstBB := copy(FirstBB);
6   CopiedBBPairs := {(FirstBB, CopiedFirstBB)};
7   **cloneBranch**(CopiedFirstBB, DescendantBBs, CopiedBBPairs);
8   NewConditionalBlock := ParentBB.splitat(Conditional);
      /* Generate the runtime check */
9   GadgetBB := createBasicBlock();
10   GadgetBB.createConditional(**all**, CopiedFirstBB, NewConditionalBlock);
      /* Insert the runtime check */
11   ParentBB.setSuccessor(GadgetBB);
12   UpdateSSA(CopiedBBPairs);
13 **end**

Algorithm 1 presents the transformation for inserting the runtime check and the code variant via a modification of LLVM-IR in the Static Single Assignment (SSA) [2] form. The pseudo-code shown is a simplified excerpt, omitting some details like the maintenance of metadata such as loop information, branch probability and vector shapes. The algorithm is applied to the branches of every IF-statement with a warp-coherent condition detected by the method described in Section 3. The input is the branch instruction *Conditional* that contains the condition and the labels of two branches of the condition, and *treatedBranchID* indicating which branch is treated in this transformation. Depending on the analysis phase, this treated branch could be either the THEN or the ELSE branch. For simplicity, we assume the THEN branch to be used without loss of generality.

At the beginning of Algorithm 1, *FirstBB* of THEN branch is located (line 1). Because of the modifying nature of our algorithm, it must not be applied to branches that can be accessed without accessing the condition first. For that reason, it has to be ensured that *ParentBB*, which contains *Conditional*, always dominates *FirstBB* of the THEN branch (line 2-3). Block B is dominated by block A iff no path to B bypassing A exists in the control flow. This prevents any interference with unrelated code that might contain calls to

modified blocks. Next, the algorithm collects all *DescendantBBs* of *FirstBB* (line 4). *DescendantBBs* consequently constitute the subtree that represents the THEN branch. For the code variant, it is now necessary to clone the THEN branch by copying each of its *DescendantBBs* (line 5-7). This is achieved by invoking Algorithm 2 described below. Afterwards, the generation and insertion of the runtime check takes place. The runtime check ought to be inserted before *Conditional*, and it is thus necessary to split *ParentBB* (line 8) at *Conditional*, resulting in two parts: *ParentBB* and *NewConditionalBlock*. We then create a new *GadgetBB* (line 9-10) that contains the *all* mask checks for selecting the optimized code variant if the corresponding mask is all-true for the elements in the current vector register (line 10). We insert *GadgetBB* between the split *ParentBB* and *NewConditionalBlock* (line 10-11). Finally, we update the structure of the control flow graph using the SSA form (line 12).

Note that partial linearization will regard *GadgetBB* that contains the *all* mask as a uniform branch, so it will retain it without linearization. The *all* mask checks if the mask is all-true, i.e., each element of a mask vector is true. It is implemented as an *all-true* mask intrinsic in WCCV. On an x86 processor with 256-bit AVX vector instructions, the IR shown in Listing 3 is generated for the *all* mask intrinsic in the LLVM-IR level and the resulting assembly is an efficient *ptest* instruction.

```
1  %cond_not = xor <8 x i1> %cond, <i1 true, i1 true, i1 true, i1
     ↪ true, i1 true, i1 true, i1 true, i1 true>
2  %ptest_zext = sext <8 x i1> %cond_not to <8 x i32>
3  %ptest_bc = bitcast <8 x i32> %ptest_zext to i256
4  %ptest_comp = icmp ne i256 %ptest_bc, 0
```

Listing 3: LLVM-IR generated for *all-true* mask intrinsic

Algorithm 2: cloneBranch

Input: CopiedFirstBB, DescendantBBs, CopiedBBPairs
Output: CopiedBBPairs

1 **for** *each* BB ∈ *CopiedFirstBB.getSuccessor()* **do**
       /* Check if BB has already been copied */
2   **if** *BB* ∈ *(DescendantBBs ∪ CopiedBBPairs.Seconds)* **then**
3   **if** *BB* ∈ *CopiedBBPairs.Seconds* **then**
               /* BB is already copied: use the copy */
4   CopiedBB := BB;
5   **else**
               /* BB still has to be copied */
6   CopiedBB := copy(BB);
7   CopiedBBPairs.add((BB,CopiedBB));
8   **cloneBranch**(CopiedBB, DescendantBBs,
                   CopiedBBPairs);
9   **end**
10   CopiedFirstBB.setSuccessor(CopiedBB);
11   RemapInstruction();
12   **end**
13 **end**

Algorithm 2 is responsible for cloning the THEN branch by recursively traversing all *DescendantBBs* that are individually copied and chained in the correct order. To achieve this, Algorithm 2 requires three input arguments: *CopiedFirstBB* representing the basic

block that is currently copied, *DescendantBBs* to be copied, and the already copied list of pairs *CopiedBBPairs*.

At the beginning, the list of *CopiedBBPairs* is initialized with *FirstBB* and its corresponding *CopiedFirstBB*. The algorithm then locates each succeeding *Block* of *CopiedFirstBB* (line 1). For each succeeding *BB* of a branch, the algorithm has to determine whether the *BB* is already copied or to copy it (line 2), followed by using the existing copy (line 4) in case of an already copied *BB* (line 3) or cloning the respective branch by recursively copying *BB* together with all its successors (line 7 -9). The maintenance of the *CopiedBBPairs* list ensures that the algorithm can handle cyclic control flow graphs without creating multiple redundant copies of *BB*. The algorithm then replaces *BB* with the newly created *CopiedBB* as a successor of *CopiedFirstBB* (line 10).

5 IMPLEMENTATION AND EXPERIMENTS

We have implemented WCCV [1] on top of Region Vectorizer (RV) [16], which is a unified vectorizer with the capability of vectorizing both loops and functions on the level of LLVM-IR. We use it to vectorize programs as an extension of the PACXX framework [8], which offers an LLVM-based, single-source, single-compiler SPMD programming model for explicit data parallelism in C++.
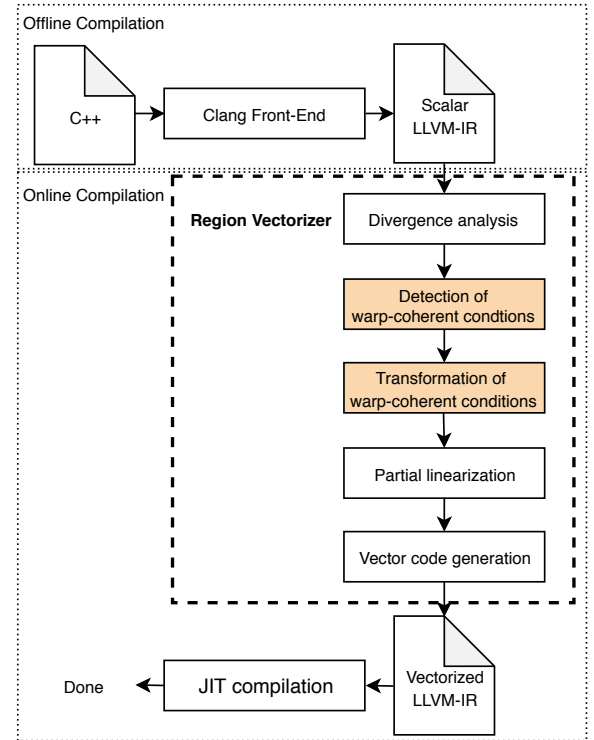

Figure 4: Implementation of WCCV on top of Region Vectorizer (RV) in the PACXX framework

Fig. 4 illustrates the PACXX compilation toolchain that is structured in two phases as follows.

---

[1]https://github.com/HuihuiSun/WCCV

a) *Offline compilation* employs a modified Clang Front-End to transform C++ code to scalar LLVM-IR.

b) *Online compilation* employs RV (within the dashed box in the figure) to transform the scalar LLVM-IR to vectorized LLVM-IR, which is then compiled and executed using the JIT capabilities of the LLVM framework. The original vectorization process of RV within the online compilation phase is subdivided into the following three passes.

b.1) *Divergence analysis* assigns a vector shape to each variable and condition: this vector shape can be either uniform (the same value for different threads) or varying (different values for different threads).

b.2) *Partial linearization* [17] uses the assigned vector shapes to partially convert control flow into data flow, while branches with conditions of uniform vector shape are preserved.

b.3) *Vector code generation* emits the final vectorized LLVM-IR that is used for JIT compilation afterwards.

WCCV is implemented by extending the compilation toolchain of RV used in PACXX with two additional passes (highlighted) between the divergence analysis and partial linearization phases, i.e., between *b.1* and *b.2*.

b.4) *Detection of warp-coherent conditions* determines the warp-coherence of each varying condition. In this pass, we implement the affine analysis, followed by the heuristic based on branch probability. We also use ATF to obtain good threshold values for the parameters in our heuristic as described in Section 3.

b.5) *Transformation of warp-coherent conditions* augments each detected warp-coherent condition with a runtime check (this check condition is uniform) and a corresponding code variant. We implement Algorithm 1 as described in Section 4 in this pass.

The input to the newly introduced passes is the annotated LLVM-IR in the SSA form, with each condition equipped with a vector shape that may be either uniform or varying. The output is the LLVM-IR in SSA, with possibly more uniform conditions that will be preserved by the partial linearization pass, thus reducing the number of masked instructions and redundant computations.

6 EXPERIMENTAL EVALUATION

We conduct a set of experiments to evaluate our WCCV approach, using the Rodinia benchmark suite and a kernel from AOBench. The benchmarks are slightly preprocessed to fit the PACXX framework, leaving the semantics unchanged.

6.1 Experiment Setup

Our evaluation platforms are listed in Table 2. We use two Intel processors–*Sandy Bridge* and *Skylake*, covering two mainstream vector instruction sets on modern computers. Each core of the processor comes with one SIMD extension. We show experimental results on one core of each processor. Each of the AVX extensions used on Sandy Bridge has 16 256-bit vector registers in x86-64 mode, while each of the AVX512 extensions on Skylake has 32 512-bit vector registers in x86-64 mode. We choose these two processors to evaluate the impact of SIMD vector width on the efficacy of our WCCV approach. The operating system is Ubuntu 16.04.5 LTS for both platforms.

Table 2: Evaluation platforms

Platform	CPU	SIMD
Sandy Bridge	Xeon E5-1620 v2 (3.70GHz)	AVX (256-bit)
Skylake	Xeon Gold 6148(2.40GHz)	AVX512 (512-bit)

In the next sections, we report runtime performance for all benchmarks in different configurations as follows:

- **RV**: performs partial control flow linearization with BOSCC gadget as in [17].
- **WCCV**: augments RV with our detection and transformation of warp-coherent conditions.

For each configuration of each benchmark, we report the mean (of 30 individual runs) speedup over the scalar version, i.e., the version without vectorization.

6.2 Performance on Sandy Bridge

Fig. 5 shows speedups achieved by WCCV and RV on Sandy Bridge. By reducing the number of masked instructions, WCCV yields more than 5% improvement over RV for 11 benchmarks, with eight of
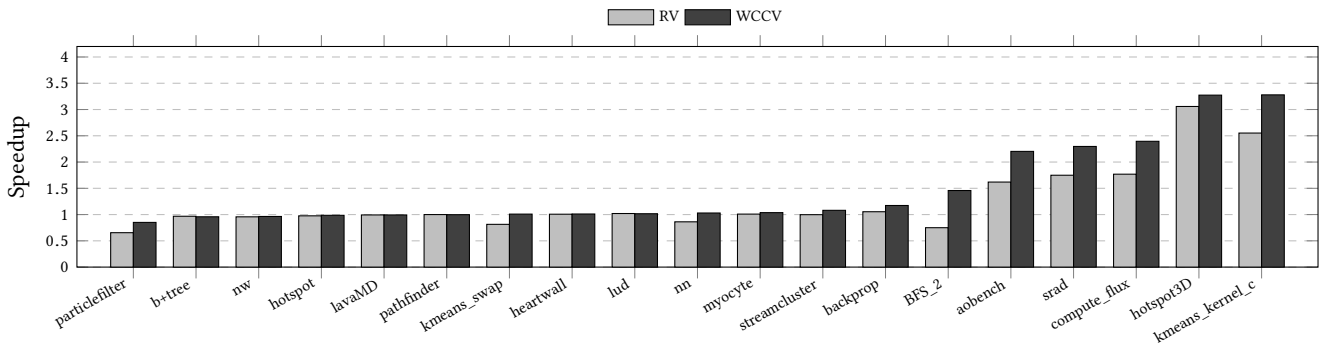

Figure 5: Runtime speedups of RV and WCCV over scalar code on Sandy Bridge

them benefiting from our technique by over 10%. Our technique leads to a mean speedup of 1.17× over RV and 1.47× over scalar code.

WCCV shows significant performance improvement over RV for 11 benchmarks, including *hotspot3D*, *streamcluster*, *backprop*, *nn*, *aobench*, *kmeans_swap*, *kmeans_kernel_c*, *particlefilter*, *srad*, *BFS_2* and *compute_flux*. These benchmarks all have a reasonable fraction of warp-coherent conditions in the code. The execution of these benchmarks is likely to branch into the optimized code variant with a high probability, reducing the number of masked instructions to a great extent and therefore improving the performance.

For comparison, RV introduces the BOSCC gadget to 7 of the 11 benchmarks, including *hotspot*, *streamcluster*, *nw*, *lud*, *srad*, *aobench* and *compute_flux*. Unlike our technique, which adds a high-performance code variant, the BOSCC gadget tries to bypass those branches that may be rarely executed during runtime.

For example, RV can speed up *compute_flux* by 1.77× over the scalar version, but WCCV obtains a speedup of 2.40×. In addition, our technique also brings performance improvement over the scalar code for benchmarks that suffer from performance degradation when vectorized with RV: *nn*, *kmeans_swap*, and *BFS_2* .

There are 7 benchmarks achieving more than 10% improvement over scalar code, with 4 of them at least twice as fast. These 7 benchmarks are suited to vectorization with continuous memory accesses. There is also one benchmark, *particlefilter*, suffering from performance degradation by 5% over the scalar version due to its random memory access pattern.

The remaining 6 benchmarks experience slight changes in performance after applying our transformations, partially offset by the overhead of barrier synchronizations. A closer investigation of the generated IR and assembly code of these benchmarks reveals that our transformation has indeed been applied, because the *all* mask intrinsics as shown in Listing 3 are found in the LLVM-IR, and ptest instructions are present in the resulting assembly code of WCCV, while this is not the case for RV. However, the performance improvement of WCCV only affects a minor fraction of the overall execution time, leading to invisible performance improvements in these cases. We anticipate performance improvements after we improve our technique with barrier synchronizations, which is

planned as future work. *Myocyte* is an exceptional case: data dependencies between values at different time steps forces only one active thread, and it may not benefit from WCCV.

No warp-coherent conditions were found in *nw* and *b+tree*, making these two benchmarks outside the scope of our technique. In addition, *nw* is dominated by barrier synchronizations.

6.3 Performance on Skylake

Fig. 6 shows speedups achieved by using RV and WCCV on Skylake. Eight benchmarks obtain a speedup of over 5%, with 6 exceeding 10%. The mean speedup achieved by WCCV is 1.14× over RV and 2.11× over the scalar version. While the overall speedup over scalar code on Skylake is better than that on Sandy Bridge, the performance improvement over RV on Skylake falls behind that on the latter. We attribute this difference mainly to the width of the SIMD vector registers used on these two platforms, as we discuss later in detail in Section 6.4.

The intriguing point is that 3 benchmarks – *hotspot3D*, *BFS_2*, and *srad* – achieve a remarkable performance improvement due to WCCV on Sandy Bridge but not on Skylake. While hotspot3D and srad still outperform RV on Skylake slightly, WCCV performs identically with RV on *BFS_2* by speeding up the benchmark by 2.47 × over scalar code. The reason is that our WCCV does not apply the transformations proposed in the paper to *BFS_2* due to the ATF-assisted heuristic, so the performance does not fall behind that of RV. In fact, applying our transformation to *BFS_2* if ATF does not suppress it would cause a 34% performance degradation, since the overhead of runtime checks is greater than the benefit of fewer masked instructions. Due to a longer SIMD vector register width than that on Sandy Bridge, the Skylake implementation has a higher coherence threshold for transformation, resulting in more masked instructions on Skylake. No benchmarks with a performance degradation of more than 5% over the scalar version are measured on Skylake.

As on the Sandy Bridge platform, 5 benchmarks including *hotspot*, *lavaMD*, *pathfinder*, *heartwall* and *lud* do not experience significant performance improvements due to WCCV because of barrier synchronization bottlenecks. For *nw* and *b+tree*, no warp-coherent conditions were found, such that the proposed transformations were not applied. The *myocyte* benchmark is indeed vectorized,
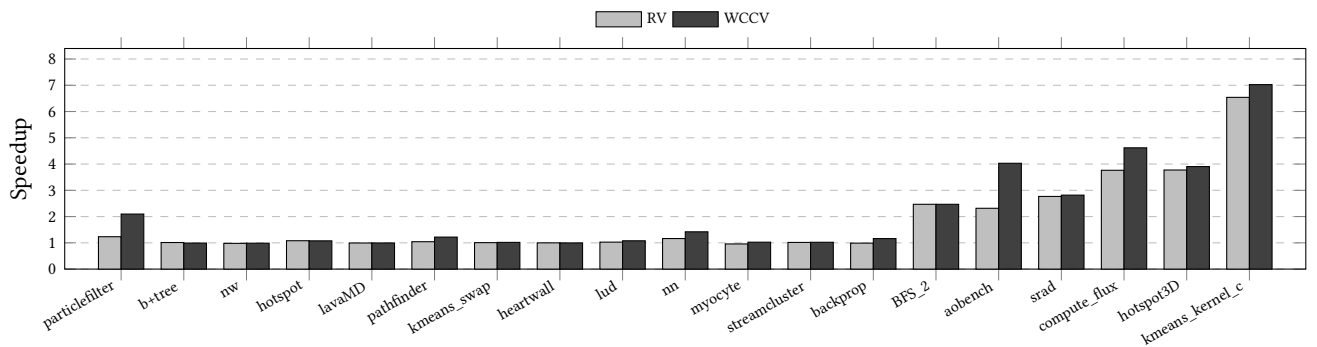

Figure 6: Runtime speedups of RV and WCCV over scalar code on Skylake

Table 3: Impact of SIMD Vector Width (Sandy Bridge/Skylake)

Kernel	Speedup over RV	Speedup over scalar	Condition type	Masked branch of RV	Masked branch of WCCV
nn	1.19/1.22	1.03/1.42	affine	42,764	4/12
kmeans_swap	1.22/1.08	1.00/1.11	affine	494,020	4/4
particlefilter	1.29/1.70	0.85/2.09	affine	40,000	0/0
srad	1.31/1.02	2.29/2.82	affine	229,916	0/0
kmeans_kernel_c	1.28/1.07	3.28/7.02	affine	494,020	4/4
aobench	1.36/1.74	2.20/4.03	high-probability	4,179,522	10,342/23,762
compute_flux	1.36/1.23	2.40/4.62	high-probability	2,147,483,600	117,767,998/176,711,998

although there is no speedup because of the small amount of parallelism.

6.4 Impact of SIMD Vector Width

The width of SIMD vector registers has significantly increased recently, resulting in more threads executing lock-step within the same vector register. As for our target platforms, Skylake has longer vector registers than Sandy Bridge. This may result in the runtime checks for WCCV passing less frequently. In this section, we examine experimentally the impact of SIMD vector width on WCCV.

Table 3 summarizes the speedup achieved by WCCV for the 7 benchmarks with performance improvement on Sandy Bridge greater than 10% over RV (except *BFS_2*, where no warp-coherent conditions are detected on Skylake). We inspect the runtime frequency of the masked branches by means of instrumented instructions, with the results also shown in Table 3.

We observe that WCCV achieves higher speedup over scalar code on Skylake, as Skylake's wider SIMD vector registers provide greater potential for vectorization improvements (also available to RV) than on Sandy Bridge. The last two columns of Table 3 show that the number of masked branches removed on Skylake is smaller. This is because our transformation is applicable less often, suppressed either during compilation by the ATF-assisted heuristic (including *BFS_2*) or during execution by the runtime check (including *compute_flux*). Accordingly, WCCV usually achieves a smaller fraction of the potential improvement on Skylake relative to RV than it does on Sandy Bridge. Our results are consistent with the experimental conclusions in [24] that increasing SIMD widths

reduces the profitability for runtime code variants, such that fewer instructions can benefit from dynamic execution paths.

6.5 Comparison with OpenCL on Sandy Bridge

In this paper, we use the PACXX framework to program in the SPMD style with C++. OpenCL is another popular SPMD programming model widely used in the parallel programming community. Fig. 7 shows experiments to compare the performance of the PACXX implementation of WCCV to Intel's OpenCL implementation on Sandy Bridge, with the speedup normalized to the fastest one in each case. PACXX uses TBB for multi-threading and RV for vectorization, shown as *RV* in the figure, followed by the version which integrates WCCV shown as *WCCV*. *OpenCL* refers to Intel's OpenCL implementation. The OpenCL result for b+tree is missing due to a runtime error in the Intel implementation.

Fig. 7 confirms that WCCV outperforms the Intel OpenCL implementation on 9 out of the 19 benchmarks, including *hotspot3D*, *particlefilter*, *lud*, *myocyte*, *BFS_2*, *aobench*, *kmeans_swap*, *kmeans_kernel_c* and *compute_flux*. The baseline PACXX framework, i.e., the RV version, is already better than the Intel OpenCL implementation for benchmarks *particlefilter*, *lud*, *myocyte* and *kmeans_kernel_c*, and it is further improved by our technique. For *hotspot3D*, *BFS_2*, *aobench*, *kmeans_swap* and *compue_flux*, the RV version of PACXX falls behind the Intel OpenCL implementation but surpasses the latter when used together with WCCV. For the remaining benchmarks, WCCV meets or exceeds RV performance, although both fall short of OpenCL, suggesting a potential for future improvement in PACXX/RV.
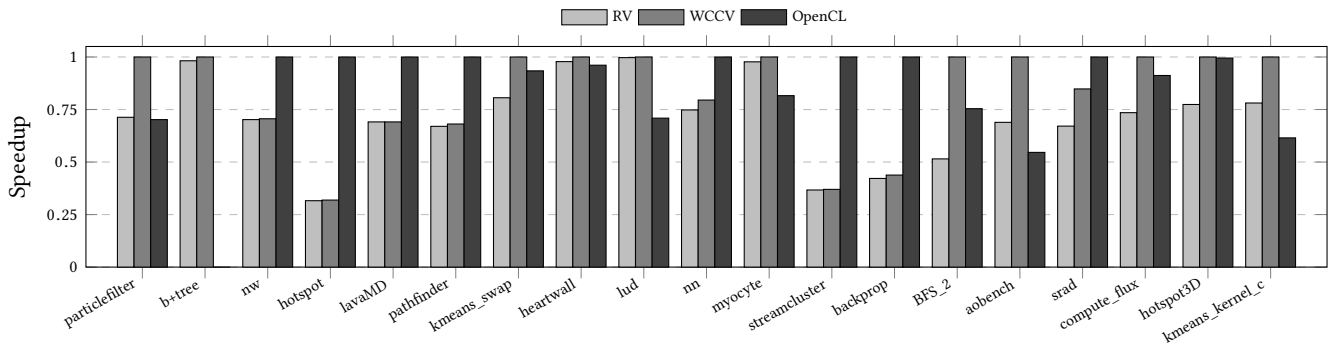
Figure 7: Relative speedups compared with Intel implementation on Sandy Bridge

For the cases where PACXX falls behind Intel OpenCL, there are mainly two reasons. First, the barrier synchronization pass of the PACXX framework splits one kernel into multiple phases at the barrier points, resulting in poor data locality for benchmarks *nw*, *hotspot*, *lavaMD*, *backprop*, *srad* and *pathfinder*. Second, the PACXX framework tries to copy data for SIMD extensions, leading to more memory accesses for benchmarks like *Streamcluster*.

7 RELATED WORK

Two approaches to vectorization are widely used in modern compilers: 1) loop vectorization [27] combines multiple occurrences of a scalar operation across consecutive loop iterations into one SIMD instruction, 2) basic block or SLP (Superword Level Parallelism) vectorization [13] replaces a group of isomorphic operations by a SIMD instruction. With the increasing popularity of SPMD programming for GPUs (Graphics Processing Units), efforts have been made to support the data-parallel programming model for CPUs by transforming multiple program instances (kernels) into one SIMD instruction to fully exploit the SIMD parallelism [5, 8, 19]. Recently, an interesting trend in parallel programming is to embed SPMD programming into C++, like in SYCL [22] and PACXX [8]. In this paper, we focus on SPMD programming with explicit vectorization capability. Our WCCV can be easily combined with the traditional implicit vectorization methods.

The idea of IF-conversion is well established in vectorizing code with IF-statements. Allen et al. [1] pioneer the work of conversing control flow to data flow by targeting vector computers with explicit hardware support for predicated execution. Shin et al. [25] extend the conventional SLP method to work in the presence of IF-statements. Sun et al. [28] present an approach to vectorizing loops with nested IF-statements for SIMD extensions without masked instructions. Pohl et al. [20] propose a set of solutions to generate efficient vector instructions in the presence of control flow for ARM NEON. Karrenberg et al. [10] present an IF-conversion approach, which is further improved by Moll et al. [17] for partial control flow linearization. Our WCCV approach uses partial control flow linearization for uniform conditions and it exploits runtime condition coherence in varying conditions to alleviate the overhead introduced by the redundant computations in the existing methods.

Divergence analysis [23] focuses on the memory divergence, improving the memory access via uniform and affine vectors. Our work follows the affine analysis used in divergence analysis, and we extend its applicability to detect boolean-step conditions. The classic loop unswitching [29] approach checks whether the condition is uniform across the loop iterations. Vectorization analysis in [10] assumes that branches with uniform conditions are never converted from control flow to data flow, and introduces branches after linearization (control flow to data flow transformation) to skip a basic block or an entire control flow path if the condition is always true or false for a single vector [17]. Such techniques make progress by resorting to static analysis, but they do not analyze the condition coherence in varying conditions.

The Branch-On-Superword-Condition-Codes (BOSCC) technique [25, 26] works by checking the emptiness of the generated masked instructions. Such vector instructions are considered redundant and are eliminated if the emptiness of their conditions is proved. The RV

technique [17] introduces a BOSCC gadget to the control flow graph after partial linearization. Our WCCV is inspired by these methods but follows a different way. In particular, the existing methods always try to pass over least frequent branches, while WCCV substitutes most frequent branches with code variants without masked instructions. Timnat et al. [30] propose duplicating code regions and employing runtime tests that check if control must pass to the linearized version or can continue to execute the uniform version. Kim et al. [12] address divergence through a multi-tier vectorization approach based on runtime convergence check. However, these methods may lead to performance degradation as they are applied to every condition. Our method provides an ATF-assisted heuristic to prevent performance degradation.

The nvcc compiler [18] inserts code to check if all threads within a warp fall in the same branch and then branch accordingly on a GPU, while we target CPUs with SIMD extensions. The ISPC compiler [19] requires programmers to express coherent control flow statements in an explicit manner to indicate performing additional tests. WCCV uses automatic compiler techniques instead for identifying coherent conditions.

8 CONCLUSIONS

In the paper, we introduce a new notion of warp-coherence of IF-conditions that have coherent behavior within the lanes of each single vector register. We present WCCV, an approach to detect and effectively exploit warp-coherence to reduce the number of redundant computations and masked instructions introduced by existing vectorization methods. WCCV is based on applying two analyses: 1) checking that the condition is an affine function, and 2) predicting the branch probability for the IF-statement. We use auto-tuning to find the optimal correspondence of benefit vs. overhead of our program transformation. We implement WCCV in an LLVM-based vectorizing compiler, and we evaluate it with the Rodinia benchmark suite on two modern processors with different SIMD extensions. Our evaluation shows significant performance improvements over existing techniques implemented in modern compilers.

REFERENCES

[1] John R. Allen, Ken Kennedy, Carrie Porterfield, et al. 1983. Conversion of Control Dependence to Data Dependence. In *Proceedings of the Symposium on Principles of Programming Languages (POPL), Austin, Texas, USA*. 177–189. https://doi.org/10.1145/567067.567085

[2] Bowen Alpern, Mark N Wegman, and F Kenneth Zadeck. 1988. Detecting equality of variables in programs. In *Proceedings of the 15th ACM SIGPLAN-SIGACT symposium on Principles of programming languages (POPL)*. ACM, 1–11.

[3] AOBench. 2019. http://code.google.com/p/aobench. accessed May 2019.

[4] Krste Asanovic, Ras Bodik, Bryan Christopher Catanzaro, Joseph James Gebis, Parry Husbands, Kurt Keutzer, David A Patterson, William Lester Plishker, John Shalf, Samuel Webb Williams, et al. 2006. *The landscape of parallel computing research: A view from Berkeley*. Technical Report. Technical Report UCB/EECS-2006-183, EECS Department, University of California, Berkeley.

[5] OpenMP Architecture Review Board. [n. d.]. *OpenMP Application Programming Interface*. accessed May 2019.

[6] Shuai Che, Michael Boyer, Jiayuan Meng, David Tarjan, Jeremy W. Sheaffer, Sang-Ha Lee, and Kevin Skadron. 2009. Rodinia: A benchmark suite for heterogeneous computing. In *Proceedings of the 2009 IEEE International Symposium on Workload Characterization (IISWC), October 4-6, 2009, Austin, TX, USA*. 44–54. https://doi.org/10.1109/IISWC.2009.5306797

[7] Free Software Foundation. [n. d.]. Using the GNU Compiler Collection (GCC). https://gcc.gnu.org/onlinedocs/gcc/. accessed May 2019.

[8] Michael Haidl, Simon Moll, Lars Klein, Huihui Sun, Sebastian Hack, and Sergei Gorlatch. 2017. PACXXv2 + RV: An LLVM-based Portable High-Performance Programming Model. In *Proceedings of the Fourth Workshop on the LLVM Compiler Infrastructure in HPC (LLVM-HPC@SC)*. 7:1–7:12. http://doi.acm.org/10.1145/3148173.3148185

[9] Intel. [n. d.]. *Intel 64 and IA-32 Architectures Optimization Reference Manual*. accessed May 2019.

[10] Ralf Karrenberg and Sebastian Hack. 2011. Whole-function vectorization. In *Proceedings of the International Symposium on Code Generation and Optimization (CGO), Chamonix, France*. 141–150. https://doi.org/10.1109/CGO.2011.5764682

[11] Khronos Group. [n. d.]. *The open standard for parallel programming of heterogeneous systems*. accessed May 2019.

[12] Hee-Seok Kim, Izzat El Hajj, John A Stratton, and Wen-Mei W Hwu. 2014. Multi-tier Dynamic Vectorization for Translating GPU Optimizations into CPU Performance. *Center for Reliable and High-Performance Computing* (2014).

[13] Samuel Larsen and Saman P. Amarasinghe. 2000. Exploiting superword level parallelism with multimedia instruction sets. In *Proceedings of the Conference on Programming Language Design and Implementation (PLDI), Vancouver, Britith Columbia, Canada*. 145–156. https://doi.org/10.1145/358438.349320

[14] Chris Lattner and Vikram S. Adve. 2004. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *Proceedings of the International Symposium on Code Generation and Optimization (CGO), San Jose, CA, USA*. IEEE, 75–88. https://doi.org/10.1109/cgo.2004.1281665

[15] Yunsup Lee, Ronny Krashinsky, Vinod Grover, Stephen W. Keckler, and Krste Asanovic. 2013. Convergence and scalarization for data-parallel architectures. In *Proceedings of the 2013 IEEE/ACM International Symposium on Code Generation and Optimization (CGO), February 23-27, 2013, Shenzhen, China*. 32:1–32:11. https://doi.org/10.1109/CGO.2013.6494995

[16] Simon Moll. [n. d.]. The Region Vectorizer(RV). https://github.com/cdl-saarland/rv. accessed May 2019.

[17] Simon Moll and Sebastian Hack. 2018. Partial Control-Flow Linearization. In *Proceedings of the Conference on Programming Language Design and Implementation (PLDI), NewYork, NY, USA*. https://doi.org/10.1145/3192366.3192413

[18] NVIDIA. [n. d.]. *CUDA COMPILER DRIVER NVCC*. accessed May 2019.

[19] Matt Pharr and William R Mark. 2012. ispc: A SPMD compiler for high-performance CPU programming. In *Innovative Parallel Computing (InPar), 2012*. IEEE, 1–13. https://doi.org/10.1109/InPar.2012.6339601

[20] Angela Pohl, Biagio Cosenza, and Ben H. H. Juurlink. 2018. Control Flow Vectorization for ARM NEON. In *Proceedings of the 21st International Workshop on Software and Compilers for Embedded Systems (SCOPES), May 28-30, 2018, Sankt Goar, Germany,*. 66–75. https://doi.org/10.1145/3207719.3207721

[21] Ari Rasch and Sergei Gorlatch. 2018. ATF: A generic directive-based auto-tuning framework. *Concurrency and Computation: Practice and Experience* (2018). https://doi.org/10.1002/cpe.4423

[22] Ruyman Reyes and Victor Lomüller. 2015. SYCL: Single-source C++ accelerator programming. In *Proceedings of the International Conference on Parallel Computing (ParCo), 1-4 September 2015, Edinburgh, Scotland, UK*. 673–682. https://doi.org/10.3233/978-1-61499-621-7-673

[23] Diogo Sampaio, Rafael Martins de Souza, Sylvain Collange, and Fernando Magno Quintão Pereira. 2013. Divergence analysis. *ACM Transactions on Programming Language and Systems* 35, 4 (2013), 13:1–13:36. https://doi.org/10.1145/2523815

[24] Thomas Schaub, Simon Moll, Ralf Karrenberg, and Sebastian Hack. 2015. The Impact of the SIMD Width on Control-Flow and Memory Divergence. *ACM Transactions on Architecture and Code Optimization* 11 (01 2015), 1–25. https://doi.org/10.1145/2687355

[25] Jaewook Shin, Mary W. Hall, and Jacqueline Chame. 2005. Superword-Level Parallelism in the Presence of Control Flow. In *3nd IEEE / ACM International Symposium on Code Generation and Optimization (CGO), 20-23 March 2005, San Jose, CA, USA*. 165–175. https://doi.org/10.1109/CGO.2005.33

[26] Jaewook Shin, Mary W. Hall, and Jacqueline Chame. 2009. Evaluating compiler technology for control-flow optimizations for multimedia extension architectures. *Microprocessors and Microsystems - Embedded Hardware Design* 33, 4 (2009), 235–243. https://doi.org/10.1016/j.micpro.2009.02.002

[27] N. Sreraman and R. Govindarajan. 2000. A Vectorizing Compiler for Multimedia Extensions. *International Journal of Parallel Programming* 28 (2000), 363–400. https://doi.org/10.1023/A:1007559022013

[28] Huihui Sun, Sergei Gorlatch, and Rongcai Zhao. 2018. Refactoring Loops with Nested IFs for SIMD Extensions Without Masked Instructions. In *Euro-Par 2018: Parallel Processing Workshops - Euro-Par 2018 International Workshops, Turin, Italy, August 27-28, 2018, Revised Selected Papers*. 769–781. https://doi.org/10.1007/978-3-030-10549-5_60

[29] J. Thomas, FE Allen, and J Cocke. 1971. *A catalogue of optimizing transformations*. Englewood Cliffs, N.J. : Prentice-Hall. 1–30 pages.

[30] Shahar Timnat, Ohad Shacham, and Ayal Zaks. 2014. Predicate vectors if you must. In *Workshop on Programming Models for SIMD/Vector Processing (WPMVP), Feburary 16th, 2014, Orlando, Florida, USA*.

[31] Haichuan Wang, Peng Wu, Ilie Gabriel Tanase, Mauricio J. Serrano, and José E. Moreira. 2014. Simple, portable and fast SIMD intrinsic programming: generic simd library. In *Proceedings of the 2014 Workshop on Programming models for SIMD/Vector processing (WPMVP), February 16, 2014, Orlando, Florida, USA,*. 9–16. https://doi.org/10.1145/2568058.2568059