Efficient Execution of Recursive Programs on Commodity Vector Hardware

Abstract

The pursuit of computational efficiency has led to the proliferation of *throughput-oriented* hardware, from GPUs to increasingly-wide vector units on commodity processors and accelerators. This hardware is designed to efficiently execute data-parallel computations in a vectorized manner. However, many algorithms are more naturally expressed as divide-and-conquer, recursive, *task-parallel* computations; in the absence of data parallelism, it seems that such algorithms are not well-suited to throughput-oriented architectures. This paper presents a set of novel code transformations that expose the data-parallelism latent in recursive, task-parallel programs. These transformations facilitate straightforward vectorization of task-parallel programs on commodity hardware. We also present scheduling policies that maintain high utilization of vector resources while limiting space usage. Across several task-parallel benchmarks, we demonstrate both efficient vector resource utilization and substantial speedup on chips using Intel's SSE4.2 vector units as well as accelerators using Intel's AVX512 units.

1. Introduction

As energy efficiency and power consumption become increasingly relevant issues for processor and accelerator designers, hardware resources for parallelism are being shifted from general-purpose multicores to *throughput-oriented* computing, with GPUs, accelerators like Intel's Xeon Phi, and increasingly-wide SIMD units on commodity processors providing efficient, vector-based parallel computation. In fact, because SIMD extensions on commodity processors tend to require relatively little extra hardware, executing a SIMD instruction is essentially "free" from a power perspective, making vectorization an attractive option.

Vector designs are well-suited to executing *data-parallel* algorithms, where the same computation is performed on each of a series of data items, and modern vectorizing compilers do a reasonable job of finding parallelism in simple, data-parallel loops and mapping that parallelism to vector units on general-purpose processors [Nuzman and Zaks 2008; Maleki et al. 2011], and programming models like CUDA and OpenCL simplify the task of mapping data-parallel computations to vector hardware on GPUs [NVIDIA; Stone et al. 2010]. Unfortunately, many algorithms are more naturally expressed as divide-and-conquer, recursive, *task-parallel* computations. Such programs do not naturally decompose into data-parallel representations—there are no dense, vectorizable loops—and hence it seems that existing vector hardware is a poor target for such programs.

```
1 void foo(int x)
2   if (isBase(x))
3     baseCase()
4   else
5     l1 = inductiveWork1(x) //l1 = x/2
6     spawn foo(l1)
7     l2 = inductiveWork2(l2) //l2 = x/2
8     spawn foo(l2)
```

(a) Simple recursive code. `spawn` creates new tasks

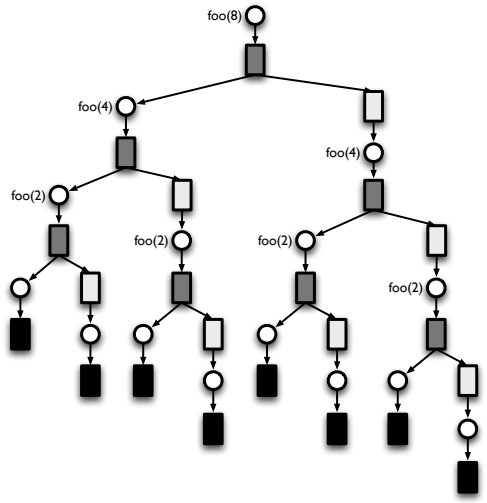

(b) Computation tree. Black boxes are `baseCase` computations, dark grey boxes are `inductiveWork1` computations and light grey boxes are `inductiveWork2` computations

Figure 1: Recursive, task-parallel code and computation tree

To address this shortcoming, there have been many proposals to map coarse-grained tasks to commodity GPUs [Tzeng et al. 2010; Aila and Laine 2009] or to modify GPU hardware to better accommodate recursive parallelism with fine-grained tasks [Orr et al. 2014; Steffen and Zambreno 2010; Huo et al. 2013]. In this paper, we consider the problem of effectively mapping fine-grained, recursive, parallel applications to *commodity vector units*. Addressing this problem would allow programmers to adopt a standard, task-parallel programming model and easily adapt existing applications to leverage the otherwise unused computational resources that exist on most general processors as well as in newer accelerators like Intel's Xeon Phi.

This paper focuses on exploiting vector parallelism on a single core. We propose code transformations that restructure recursive, task-parallel applications to expose their latent data-parallelism that allows for efficient vectorization. A typical divide-and-conquer application can be thought of as a *computation tree*, with each interior node in the computation tree representing work done prior to making a recursive call, children of a node in the tree representing the

work done during each recursive call, and leaf nodes representing work done during the base case. Figure 1 shows an abstract recursive code—the paper's running example—and its associated computation tree. An execution of the application is equivalent to a valid walk of the tree. In particular, the normal sequential execution of this computation can be represented by a depth-first walk of the tree.

Contributions: The key contributions of this paper are code transformations that create a tree-walk that can be efficiently vectorized. The transformations handle three important issues: (1) Expose data-parallel computation by performing a *breadth-first expansion* of the computation tree; (2) Reduce the amount of space used and the number of cache misses by switching to *depth-first execution* when enough parallelism has been generated; (3) When irregularities in the computation tree cause reduction in available parallelism, regenerate parallel work using *re-expansion*. In addition, we develop *block management* schemes, including a novel *stream compaction* algorithm to ensure that parallel work and data accesses remain structured for efficient SIMDization.

In our experimental evaluation, we find that our techniques can find vectorization opportunities in all the benchmarks considered, ranging from small microbenchmarks to larger kernels. On two hardware platforms, an Intel Xeon E5 with the SSE4.2 instruction set and an Intel Xeon Phi with the AVX512 instruction set, we obtain up to 12.23× speedup. We further find that our scheduling policy is effective at maintaining high SIMD utilization while bounding space usage and incurring relatively low overheads. Overall, this paper presents the first set of techniques for mapping general, recursive task-parallel programs to commodity vector hardware. Our approach allows programmers to leverage the "free" execution resources available in SIMD units even for programs that do not appear to be amenable to data-parallel vectorization.

2. Preliminaries

Specifying recursive, task-parallel programs: This paper targets the vector parallelization of recursive, task-parallel applications. To clarify the types of applications we transform and parallelize, we consider a language for specifying recursive, task-parallel programs, as defined in Figure 2. The language is a variant on Cilk [Blumofe et al. 1995; Frigo et al. 1998]. We emphasize that this language is meant to clarify the types of programs we tackle; in our implementations, we transform and evaluate programs written in C that conform to the restrictions of this language.

A k-ary recursive method evaluates a conditional (which may be a boolean expression or a function returning a boolean) to decide whether to execute the base case or the inductive case. The base case is used to produce computation results. Base case statements can assign expression results to local variables (note that expressions can be arbitrary, stateless, non-recursive functions) or perform *reductions* over one of a set of reducer objects [Frigo et al. 2009]—these

$$
\begin{aligned}
v &\in \mathbb{Z} && [Values]\\
a &\in \{a_1, a_2, \ldots a_k\} && [Arguments]\\
l &\in \{l_1, l_2, \ldots\} && [Locals]\\
r &\in \{r_1, r_2, \ldots\} && [Reducers]\\
\odot &::= < \mid > \mid = \mid \neq \mid \geq \mid \leq\\
b \in BExprs &::= \textbf{true} \mid \textbf{false} \mid e \odot e \mid f_b(e_1, e_2, \ldots)\\
e \in Exprs &::= v \mid l \mid a \mid f_v(e_1, e_2, \ldots)\\
s_b \in BaseStmts &::= s_b; s_b \mid l := e \mid \mathsf{reduce}(r, e)\\
s_i \in IndStmts &::= \textbf{return} \mid s_i; s_i \mid l := e \mid \mathsf{spawn}(e_1, e_2, \ldots, e_k)\\
m \in MethodBody &::= \textbf{if } b \textbf{ then } s_b \textbf{ else } s_i
\end{aligned}
$$

Figure 2: Language for recursive, task-parallel methods

associative, commutative updates to global state are used in lieu of return values. Note that this means that the execution of multiple base-case tasks can be readily parallelized. While the use of reduction objects instead of return values may seem limiting, we have found that many recursive applications can be written in this manner.

The inductive case can perform additional computations and make recursive calls using the spawn command, which binds expression values to the arguments of the subsequent recursive invocation. As in Cilk, spawned methods can be executed in parallel with (and are assumed to be independent of) any subsequent work in the spawning method; this is the source of task parallelism in our language.[1]

There is an implicit synchronization at the end of each method: all spawned (callee) methods must return before their parent (caller) method can return. Unlike in Cilk, our language does not have an explicit `sync` keyword. No additional work can be performed after spawned tasks "rejoin" execution. All computations expressed in our language can be viewed as computation trees: `spawns` create children of the current task, and base case computations, which perform no spawns, are leaves of the computation tree.

In terms of our language description, Figure 1(a) can be interpreted as follows. `foo` defines the recursive method. `isBase()` performs some computation to decide whether to perform the base case, which is defined by `baseCase()`. If `isBase()` returns false, `inductiveWork1()` and `inductiveWork2()` perform the necessary computations to set up two `spawns` of recursive tasks. While the running example only has two children tasks, in general, any number of child tasks can be spawned in the inductive case.

Strawman vectorization: To grasp the difficulties involved in vectorizing a recursive application described in our language, it is helpful to understand why the obvious solution will not work. Consider executing a task-parallel program written in our specification language using a traditional multicore, work-stealing runtime, as used by Cilk [Blumofe et al. 1995; Frigo et al. 1998; Danaher et al. 2006]. In a

[1] We only consider self-recursive programs in this paper for simplicity; this is not a fundamental limitation of our technique.

Cilk-style work-stealing runtime, a computation tree is run in parallel using a "work-first" scheduling policy [Frigo et al. 1998], where a thread executes a computation tree depth-first — when a thread spawns a task, it immediately executes the spawned task and places the executing task's "continuation" (the remaining work of the function) in a local pool. Other threads that need work may steal continuations to execute the remainder of the computation. In the absence of work stealing (i.e., if every thread has sufficient work), this policy results in each thread executing a subtree of the computation tree in a depth-first manner.

One obvious approach to vectorization is to map this basic execution strategy to vector units: at a high level, a thread can be assigned to each SIMD lane of a vector unit, and each thread picks a node in the computation tree and executes it in a vector-parallel manner with (some) other nodes in the computation tree, and then proceeds to the next node in a depth-first manner.

Implementing this strategy on SIMD units is extremely difficult: because each "thread" executes a different portion of the computation tree, the threads' stacks grow and shrink at different times. All of this stack management must be done manually, as all the SIMD lanes are under the control of a single, actual thread, necessarily incurring extra overhead. Moreover, performing the stack management in a vector-friendly manner is impossible: because the stacks diverge, storing/loading data from each thread's stack will require scatter and gather operations, which perform poorly on vector units designed for packed loads and stores.

3. From Task Parallelism to Data Parallelism

This section overviews how a recursive, task-parallel program can be transformed to enable vector-parallel execution. Rather than implementing our schedulers as runtime components separate from the task-parallel application, as in traditional multicore implementations, our approach to vectorization uses code transformations that integrate scheduling decisions into the (transformed) application code. That is, we transform the application code to produce particular execution schedules. We choose this approach to facilitate vectorizing fine-grained tasks. The overheads of runtime scheduling are tolerable when parallelism can be achieved by threads that each run large numbers of tasks independently. However, exploiting vector hardware requires fine-grained parallelism: operations must be grouped together at the granularity of *individual instructions* to be vectorized.

The key insight behind our vectorization strategy is that through careful code transformations, recursive, task-parallel algorithms can be transformed into *blocked* recursive algorithms, which group together multiple tasks in the original computation tree into blocks that can be efficiently executed in a vectorized manner with low overhead. These transformations have two effects: (i) by building these computation blocks out of tasks in the tree that are all at the same depth, our transformations avoid the stack management pitfalls that compromise the naïve solution described previously; and (ii) by creating blocks out of individual fine-grained tasks, our transformations enable the instruction-by-instruction grouping necessary for vectorized execution.

Our vectorization strategy consists of three components:

1. We transform the original recursive, task-parallel code into blocked code that executes the computation tree *level-by-level* in breadth-first manner. Breadth-first expansion exposes opportunities for parallelism, the blocked structure of the code enables vectorization, and the level-by-level strategy ensures that the stack frames necessary for vectorized computation can be organized to support vectorized memory operations.
2. A pure breadth-first execution can consume large amounts of space (proportional to the width of the computation tree) and lead to a large number of cache misses due to decreased locality. Therefore, we produce a second transformed version of the code that implements a *blocked depth-first* execution schedule that essentially spawns "threads" for each task in a block of tasks. Each thread explores its portion of the computation tree in a depth-first manner, and the threads execute in lockstep, each taking identical paths through their respective computation subtrees. By executing in a depth-first manner, the amount of storage required for saving state is proportional to the depth of the tree, and by executing in lockstep, each "thread" is kept at the same depth of the tree as the other threads in the block, simplifying stack management.
3. Since some branches of the computation tree are shallower than others, some threads may "die out" early, reducing SIMD utilization. To ameliorate this, we design a *re-expansion* mechanism that toggles between breadth-first execution to generate more parallel work and depth-first execution to control space usage.

4. Transformations and Scheduling

This section describes the three techniques discussed in Section 3 in more detail. We focus primarily on the code transformations necessary to achieve particular scheduling policies. The details of how this transformed code can be efficiently vectorized are in Section 5.

4.1 Breadth-first execution to extract data parallelism

Our first transformation produces a *breadth-first, level-by-level* traversal of the computation tree to generate large blocks of work that can be readily vectorized. Figure 3 shows the transformed code for the code example in Figure 1(a).

The essential idea of the transformation is that each invocation of `bfs_foo` executes all of the instances of `foo` in a given level of the tree before proceeding to the next level of the tree. Each task instance is assigned to a `Thread` structure, which contains the information that would be in the stack frame for that task instance (specifically, any arguments to the task). A `ThreadBlock` contains threads for each task at a given level of the computation tree. `bfs_foo`

```
1  void bfs_foo(ThreadBlock tb)
2    ThreadBlock next
3    foreach (Thread t : tb)
4      if (isBase(t.x))
5        baseCase()
6      else
7        l1 = inductiveWork1(t.x)
8        next.add(new Thread(l1))
9        l2 = inductiveWork2(t.x)
10       next.add(new Thread(l2))
11   bfs_foo(next)
```

Figure 3: Breadth-first version of code in Figure 1(a)

is initially called with a thread block containing a single thread whose `x` field is set to the original parameter to `foo`.

The transformed code is straightforward. At each `spawn` directive, rather than invoking the next method, the code creates an additional thread for the next task, with the appropriate arguments, and places it into the `next` thread block for the next level of the computation tree. Once all of the computation at the current level of the tree has been completed, the transformed code invokes `bfs_foo` on `next`, moving to the next level of the computation tree.

This transformation has several effects. First, consider the loop in line 3 in Figure 3. This is a dense loop over a vector (of `Threads`). Through a combination of loop distribution, inlining, if-conversion, and other standard compiler transformations, this loop can be transformed into a series of dense loops over individual instructions, which can then be readily vectorized. Note that the order in which tasks at a given level are executed can change after loop distribution: for instance, all the left children of the current level can be added to the next thread block before all the right children. This reordering is (a) still compatible with the parallel semantics of our language; and (b) potentially beneficial to vectorization, as in many task-parallel applications, left children behave similarly and right children behave similarly. The most challenging task in vectorization is vectorizing the addition of new `Threads` to the `next` block in lines 8 and 10. Section 5 describes a general *stream compaction* mechanism that can manage the blocks in an efficient, vectorized manner.

The second effect of this transformation is that it quickly generates substantial amounts of parallel work. Although the initial thread block has but one thread in it, at each level, the block gets larger, creating additional parallel work. While this feature is beneficial for keeping the vector units busy and maintaining high utilization, for large computation trees, the size of these blocks can get prohibitive; the total amount of state that must be tracked can get as large as the width of the computation tree. Moreover, as the thread blocks get larger, the code begins to suffer from poor cache performance: by the time execution moves to the next level of the computation tree, the `Threads` added to the next thread block will have been evicted from cache.

4.2 Depth-first execution to limit space usage

To overcome the space explosion incurred by the breadth-first execution strategy, we make the following observation. Suppose we stop the controlled breadth-first execution after a certain level, and let each thread in the resulting thread block execute its computation subtree to completion, as in Figure 4(a). In other words, after some number of rounds of running `bfs_foo`, we invoked `dfs_foo` instead. Each thread at the level where breadth-first execution is stopped thus executes its computation subtree in a depth-first manner by invoking the original recursive code. This execution strategy *no longer increases space usage exponentially*. In particular, if there are T threads in the thread block when `dfs_foo` is invoked, and the depth of the computation tree is D, the space usage is $O(TD)$.

The downside to this execution strategy is that the loop in line 2 of Figure 4(a) is not as easily vectorizable as the dense loop in Figure 3. While the loop is still dense, traditional techniques for vectorizing dense loops do not handle recursive methods. So have we merely saved space at the expense of losing vectorization?

In recent work, Jo and Kulkarni [2011] proposed a compiler transformation called *point blocking* that targets *repeated recursive traversals of trees*. In particular, for code that performs multiple recursive traversals of a tree in parallel, point blocking transforms the code so that multiple traversal threads are *blocked* together, and the blocks of threads traverse the tree in lockstep. For applications such as Barnes-hut, when multiple traversals are performed in lockstep, each thread in the block operates on the same part of the tree structure in close succession, leading to improved locality. Jo et al. [2013] later observed that the code structure generated by point blocking made such tree traversal codes amenable to vectorization.

The key insight for our transformation is that when a block of threads traversing the computation tree each execute their subtrees to completion, they are performing *repeated recursive traversals* not of a literal tree (as in Jo and Kulkarni's work), but of an abstract *computation tree*. While each thread does not "traverse" (execute) exactly the same computation tree, they each dynamically unfold their computation tree by executing the same code. This the same as each thread traversing a single tree, but performing slightly different work at each node in the tree. Point blocking can be directly applied to the code in Figure 4(a) to produce a new, *blocked depth-first* execution where all the threads in the block execute their computation trees in lockstep.

Figure 4(b) shows the result of applying point blocking to the depth-first code. The key to the transformation is that rather than creating a single thread block for the next level of computation, a separate thread block is created for each `spawn` directive in the code. Then the depth-first version of the code is called for each thread block in succession, so every thread executes its left subtree (to completion) before executing its right subtree. Figure 4(c) shows the order of computation imposed by the transformation *after the first two levels of the computation tree are executed in a breadth-first manner*. Just as in the breadth-first code, all the threads in a thread block are at the same level of the tree. Unlike breadth-first code, the thread blocks for the next level of the

```
1  void dfs_foo(ThreadBlock tb)
2    foreach (Thread t : tb)
3      if (isBase(t.x))
4        baseCase()
5      else
6        l1 = inductiveWork1(t.x)
7        foo(l1)
8        l2 = inductiveWork2(t.x)
9        foo(l2)
```

(a) Depth-first execution after breadth-first execution

```
1  void blocked_foo(ThreadBlock tb)
2    ThreadBlock left, right
3    foreach (Thread t : tb)
4      if (isBase(t.x))
5        baseCase()
6      else
7        l1 = inductiveWork1(t.x)
8        left.add(new Thread(l1))
9        l2 = inductiveWork2(t.x)
10       right.add(new Thread(l2))
11   blocked_foo(left)
12   blocked_foo(right)
```

(b) Blocked depth-first execution

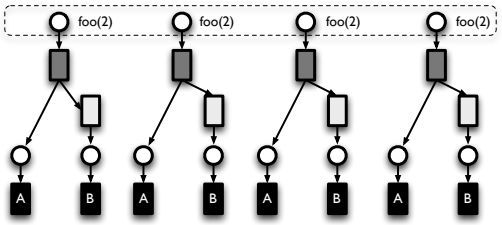

(c) Schedule of computation for blocked code after first two levels have been executed in breadth-first manner. Leaf nodes with the same label are executed as part of the same block.

Figure 4: Depth-first version and computation schedule

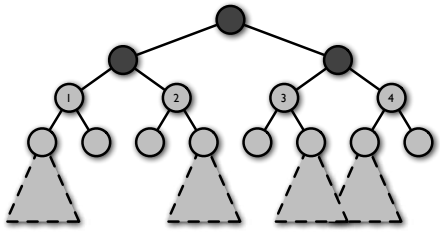

Figure 5: Computation after partial breadth-first execution

tree can have no more threads than the thread block at the current level; therefore, space usage is contained.

The transformed code can be vectorized in the same way as the breadth-first code. As in the breadth-first code, the depth-first code naturally groups together corresponding children with each other; each thread block for the next level only contains children from one `spawn` directive. Because different `spawns` in a task often behave differently, this scheduling strategy promotes similarity of tasks that are vectorized together, reducing vector divergence.

There is a downside to blocked depth-first execution: threads can only be executed in parallel if they both visit the "same" node in their computation tree (in other words, if the computation trees overlap). If one thread in a block executes its base case while the other threads continue recursing, the size of the next level block will be smaller. If a block becomes too small, there may no longer be enough threads in the block to keep all the SIMD lanes in a vector unit occupied, resulting in *underutilization* and lost parallelization opportunities. For instance, consider the stylized computation tree in Figure 5, with the dashed triangles representing the rest of the tree. If breadth-first expansion has executed the black nodes of the computation tree, there are now four threads ready to execute the grey portions of the tree. Blocked depth-first execution will cause the four threads to execute their code in lockstep. However, threads 1 and 4 in Figure 5 have left-biased computation trees, while 2 and 3 have right-biased subtrees. While threads 1 and 4 execute their left subtrees, 2 and 3 must sit idle. With only two active threads in a thread block, we cannot fully utilize even 4-way vector. The next section describes a scheduling policy that can address this underutilization.

```
1  void bfs_foo(ThreadBlock tb)
2    ThreadBlock next
3    foreach (Thread t : tb)
4      /* same as foreach in Figure 3 lines 4–10*/
5    if (next.size() < max_block_size)
6      bfs_foo(next)
7    else
8      blocked_foo(next)

10 void blocked_foo(ThreadBlock tb)
11   ThreadBlock left, right
12   foreach (Thread t : tb)
13     /* same as foreach in Figure 4(b) lines 4–10 */
14   if (left.size() > reexpansion_threshold)
15     blocked_foo(left)
16   else
17     bfs_foo(left)
18   if (right.size() > reexpansion_threshold)
19     blocked_foo(right)
20   else
21     bfs_foo(right)
```

Figure 6: Re-expansion pseudocode

4.3 Re-expansion to improve utilization

We propose a scheduling strategy called *re-expansion* to mitigate the under-utilization that can arise due to lack of overlap between different threads' computation trees. Essentially, re-expansion toggles back and forth between breadth-first execution and depth-first execution: the former to generate work when thread block sizes get too small, the latter to execute work in bounded space when thread block sizes get too large. For example, if re-expansion were applied to the computation tree of Figure 5, then after threads 2 and 3 drop out of the left portion of the depth-first computation, threads 1 and 4 can switch back to breadth-first execution, generating more work to run in parallel. Intuitively, re-expansion looks for more parallel work in the subtrees of the "live" threads during depth-first execution.

Implementing re-expansion is straightforward: because both the breadth-first code and the blocked depth-first code take thread blocks as arguments, each can call the other to switch execution strategies. Figure 6 shows how re-expansion can be integrated into the transformed code.

Re-expansion requires two thresholds, a `max_block_size` that triggers depth-first execution when the blocks are getting too big, and a `reexpansion_threshold` that triggers breadth-first execution when there is too little parallel work. These thresholds are application-specific, as they are

governed by the structure of the computation tree. To set these thresholds, we pick a target space utilization, T_{max} (i.e., the maximum number of threads we want active at a time), and determine the expansion factor, e, of an application (the maximum number of `spawn`s in a task). We set both `max_block_size` and `reexpansion_threshold` to T_{max}/e, so that after one round of breadth-first execution we cannot create more than T_{max} threads.

5. Effective SIMD Implementation

The discussion so far has focused on maximizing opportunities for vectorization by exposing the data parallelism latent in recursive-parallel programs. In this section, we discuss the mechanisms employed to translate this opportunity into actual performance. This involves replacing operations on individual threads with operations that span the entire thread block, maximizing the use of vector instructions in place of scalar instructions, and improving the data and operation structures to enable vectorized execution. We note how each aspect of a function body—stack management, base case check, base case and recursive execution—can be optimized. We present the implementation details in terms of our running example.

Optimized stack operations: Performing a blocked depth-first recursive call or a breadth-first re-expansion allows the stack operations of individual threads to be optimized. We exploit the fact that all recursive calls invoke the same function, merging the stack frames of individual threads into a thread block, which is allocated and deallocated with a constant number of instructions. The stack management overhead thus reduces with increase in block size. Within each thread block, all instances of individual data elements across all stack frames are stored contiguously. This structure-of-arrays layout lets us avoid expensive scatter/gather operations and simply replace the scalar stores and loads in individual threads with the corresponding vector instructions.

Vectorizing operations: The first operation a task performs is to check whether to execute the base case or the recursive case. This operation, denoted by `isBase()`, is performed by all threads and can be readily vectorized. The code is transformed into an iterative loop that performs the `isBase()` computation across all threads in a block. This loop is then vectorized by the compiler. In general, we use the compiler's vectorization support where possible, and introduce explicit vector instructions only where necessary. This way, we rely on the compiler to manage register allocation, scalar optimizations, and choose appropriate instruction sequences.

The result of executing `isBase()` is a vector of boolean flags (characters or bits depending on the instruction set) that denotes whether the branch is to be taken by each thread. The base and recursive cases in the different threads can now be executed using vector instructions in which elements of the vector are masked using the boolean flags. However, this would significantly complicate vector code generation. Not all scalar instructions have masked vector equivalent instructions. In addition, such masked execution significantly degrades vector utilization and thus performance.

Stream compaction: Utilization can be improved by partitioning the threads into groups that perform identical actions. All threads performing the base case need to be separated from those performing the recursive case. Once grouped, the threads performing the same action, be it base or recursive case, can be vectorized without masking. For breadth-first re-expansion, it is also beneficial to sort the recursive calls based on their *sibling position*: the number of recursive calls that precede a given invocation. The ordering of the recursive calls is ensured during breadth-first expansion by enqueuing the i-th recursive call by all threads before any $(i+1)$-th calls. Grouping the threads into those executing base case and recursive case is performed using *stream compaction*:

```
1 foreach (Thread t : tb)
2   if (t.isBase) baseCase.add(t)
3   else recursiveCase.add(t)
4 //vectorized execution of baseCase threads
5 //vectorized execution of recursiveCase threads
```

The most efficient approach to vectorizing the stream compaction operation—the foreach loop in the above code snippet—depends on the instruction set and space requirements. We now discuss how the stream compaction operation can be vectorized on the Xeon E5 (SSE4.2) and the Xeon Phi (AVX512). The Xeon E5 supports the shuffle instruction that can perform an in-place permutation of the contents of a vector register. Stream compaction corresponds to a permutation that gathers the threads taking the same branch path. This shuffle operation can be encoded as:

```
1 pos=0
2 shuffleOp = Thread[tb.size()]
3 foreach (Thread t : tb)
4   if (t.isBase) shuffleOp[pos++] = t
```

We further optimize this loop by pre-computing the `shuffleOp` value for all possible boolean vectors into a shuffle table. For a vector width—-the number of elements can be processed by a single vector instruction—t, there are 2^t possible entries in the shuffle table. Stream compaction now involves one lookup into this table to determine the desired shuffle and executing the vector shuffle instruction. While efficient in time, the space overhead of the shuffle table is exponential in the vector width. We address this by computing the shuffle to be performed using a smaller shuffle table and a multi-pass algorithm. This is conceptually similar to factorization-based implementations of various permutation operations [Eklundh 1972; Puschel et al. 2005].

Let us consider the compaction of a vector X into another vector Y, denoted by $\mathsf{compact}(X[0:N] \rightarrow Y[0:N])$ We observe that this can be factorized as:

$$\mathsf{compact}(X[0:m] \rightarrow Y[0:\mathsf{nnz}(X[0:m])]);$$
$$\mathsf{compact}(X[m+1:N] \rightarrow Y[\mathsf{nnz}(X[0:m])+1:N])$$

where $\mathsf{nnz}(X[a:b])$ is the number of predicates of interest (e.g., the number of non-zeroes) in vector X between posi-

tions a and b. In addition to the shuffle table, we pre-compute and store the nnz() function into an advance table, denoting how far the position of the next compaction needs to be advanced. Note that the table size is exponential in the vector width, while the factorized compaction requires number of instructions linear in the number of factorization steps. For example, we can reduce the size of the shuffle tables by a factor of 256 (from 2^{16} to 2^8) by using an 8-way table instead of a 16-way tables. This incurs only a few additional instructions rather than 16 that would be required by a sequential compaction. As vector width increases, as can be expected on future systems targeting energy-efficient performance improvements, the benefits from this approach improve further.

The current generation Xeon Phi does not have a vector shuffle instruction. However, it has a masked scatter operation that can store a subset of the elements in the vector into memory. We observe that the mask for the scatter operation can be computed as an exclusive prefix sum. An exclusive prefix sum of a vector X into vector Y is defined as:

$$Y[i] = \sum_{j=0}^{j<i} (X[j] \text{ should be compacted } ? \ 1 : 0)$$

As in the case of the shuffle table, we store the prefix-sum function into a table; the prefix sum computation can be factorized when combined with the advance table. Thus the space overhead can be reduced at the expense of a few additional instructions to compute the masked scatter instruction. Therefore for both Xeon E5 and Xeon Phi, we can perform stream compaction in a vectorized fashion with low space and time overhead.

6. Evaluation

We now empirically evaluate the performance of our techniques across eight recursive benchmarks. We note that vectorization of recursive benchmarks introduces overheads of various kinds. The data-parallel rather than strict depth-first execution can increase register pressure as well as the cache footprint of each function invocation. As the block size gets larger, the footprint can exceed the cache sizes, degrading cache locality. Stream compaction incurs table lookup costs, additional instructions, and memory operations that introduce additional overheads. In addition, the benefits of vectorization are limited by both the availability of enough concurrency (for instance, due to the presence of scaler instructions that are not effectively vectorized by the compiler across threads) and the ability of the blocked depth-first and breadth-first schemes to expose this concurrency in the form of data parallelism. This section shows that the vectorization gains from our techniques outweigh the overheads, across most of our benchmarks.

6.1 Evaluation platform and benchmarks

We evaluate our transformations on the Intel E5-2670 and the Intel Xeon Phi. The E5 is a 8-core 2.6GHz Sandy Bridge processor with 32KB L1 cache per core, 20MB last-level cache, and 128-bit SSE 4.2 instruction set[2]. The Xeon Phi is a 61-core SE10P co-processor running at 1.1GHz with 32KB L1 cache and 512KB L2 cache per core, supporting 512-bit AVX512 instructions. All benchmarks were compiled with Intel icc-13.3.163 compiler and `-O3`. The Xeon Phi experiments were conducted in the *native mode* with `-mmic` option. Recall that our focus is single-core vectorization: all of our experiments use a single core of the target platform.[3]

Benchmark	Problem	#levels	#tasks	Time (s)	
				E5	Phi
knapsack	long	31	2.15B	8.7	83.5
fib	45	45	3.67B	9.0	84.3
parentheses	19	37	4.85B	10.5	69.9
nqueens	13	14	59.8M	4.9	48.2
graphcol	3(38-64)	39	42.4M	30.9	417.6
uts	20	1572	136K	21.4	164.5
binomial	C(36,13)	36	4.62B	8.3	74.1
minmax	4×4	13	2.42B	18.1	120.8

Table 1: Benchmarks. All benchmarks use 16-wide vector operations, except knapsack and UTS on the Xeon E5, which employ 8-wide and 4-wide vector operations, respectively.

We evaluated our technique on the eight benchmarks, ranging from microbenchmarks to larger kernels: (i) `knapsack`, which computes the optimal solution to the knapsack problem [Cilk][4]; (ii) `fib`, which Computes the 45-th Fibonacci number [Cilk]; (iii) `parentheses`, which computes the number of well-formed parentheses string combinations with 19 parentheses; (iv) `nqueens`, which counts the number of valid solutions to the 13-queens problems [Barcelona OpenMP Task Suite (BOTS)]; (v) `graphcol`, which counts the number of valid ways of coloring a 38-node, 64-edge graph with 3 colors [Huo et al. 2013]; (vi) `uts`, which counts the number of nodes in a probabilistic binomial tree [Olivier et al. 2007]; (vii) `binomial`, which recursively computes the combination ${}_{36}C_{13}$ [Huo et al. 2013]; and (viii) `minmax`, a min-max search for tic-tac-toe on a 4×4 board.

Table 1 characterizes the benchmarks and their sequential execution time. We present speedups relative to these sequential times in the rest of the evaluation. We use the smallest data type possible without loss of generality to maximize vector width (e.g., we define n in `fib` as a `char` on E5 due to the exponential nature of the computation). On the Phi, we use the `int` data type for all benchmarks, because the IMCI instruction set does not support shorter data types well.

Figure 7 characterizes the structure of each benchmark's computation tree. Because `binomial` and `minmax` have similar trees to `fib` and `nqueens`, respectively, their char-

[2] We do not use AVX, as it does not support shuffle instructions

[3] In addition to pure SIMD execution, we have modified our transformations to be compatible with standard, Cilk-style multicore parallelism, letting other workers steal entire blocks of work during depth-first execution. This extension is not a contribution of this work, but Appendix C in the supplemental material presents preliminary results demonstrating that our SIMDization approach is complementary to multicore parallelism.

[4] We use the 'long' input without pruning to ensure determinism.

Benchmark	Xeon E5					Xeon Phi				
	Breadth-first	No Re-expansion		Re-expansion		Breadth-first	No Re-expansion		Re-expansion	
	only speedup	Block	Speedup	Block	Speedup	only speedup	Block	Speedup	Block	Speedup
knapsack	1.17	2^{12}	1.90	2^{11}	1.91	OOM	2^{8}	5.23	2^{8}	5.10
fib	1.67	2^{18*}	1.99	2^{9}	2.03	0.65	2^{10}	3.07	2^{9}	3.50
parentheses	1.23	2^{14}	1.84	2^{11}	1.85	OOM	2^{9}	1.32	2^{9}	1.39
nqueens	4.38	2^{23}	5.10	2^{15}	6.33	0.83	2^{22}	1.18	2^{12}	2.96
graphcol	1.08	2^{21}	2.99	2^{8}	8.95	0.79	2^{21}	1.88	2^{8}	12.23
uts	1.68	2^{14}	1.69	2^{14}	1.68	1.0	2^{14}	2.05	2^{14}	2.05
binomial	1.14	2^{18}	1.38	2^{18}	1.39	OOM	2^{11}	1.76	2^{9}	1.99
minmax	0.83	2^{20}	1.79	2^{10}	2.17	OOM	2^{13}	0.61†	2^{8}	0.93†
Geometric mean	1.44		2.13		2.58	0.81		1.78		2.76

* Performance is close to that for 2^9 block size † The poor performance of minmax is due to excessive cache misses in the Xeon Phi's small cache; if the cache is warmed up for the kernel computation, we can achieve a speedup of 1.09 without reexpansion and and 1.49 with (not counting the warmup).

Table 2: Best block size and execution times for different vectorization strategies

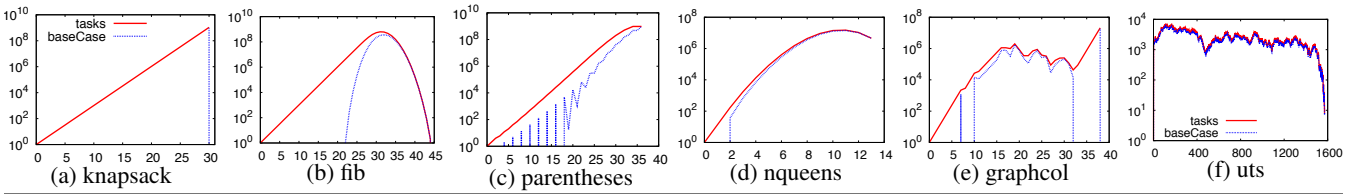

Figure 7: Distribution of tasks in selected benchmarks. x-axis: recursion depth; y-axis: number of all and base case tasks

acteristics are described in Appendix A in the supplemental material. For each benchmark, we show the number of levels, the total number of tasks in each level, and the number of tasks executing the base case in each level. `knapsack` is a perfectly balanced tree with base case tasks only at the last level. `fib` (`binomial`) and `parentheses` are more unbalanced, with `parentheses` having some intermittent shallower branches. `nqueens` (`minmax`) has a large number of leaves at almost all levels and a large fanout. `graphcol` and `uts` have a more uneven distribution of total tasks and leaves. `uts` is a deep computation tree with the fewest number of tasks in each level.

6.2 Overall speedup from blocked SIMD execution

Table 2 shows the overall speedup of our vectorized execution strategies on the E5 and Xeon Phi architectures. Pure breadth-first execution sometimes runs out of memory on Xeon Phi, and provides poor performance in general. This is due to the fact that it has poor cache performance due to large block sizes. With our hybrid depth-first/breadth-first strategy, without re-expansion, we achieve speedups of 1.38–5.10× (geometric mean of 2.13×) on the E5, and a 0.59–5.23× (geometric mean of 1.78×) on the Xeon Phi. Adding re-expansion raises speedups to 1.39–8.95× (geometric mean of 2.58×) on the E5 and 0.93–12.23× (geometric mean of 2.76×) on the Xeon Phi. Using re-expansion typically uses less space since it yields equivalent or better speedups at smaller block sizes.

6.3 Understanding vectorized performance

We now explore the various factors that affect vectorized performance in detail. As mentioned above, `binomial` and `minmax` are structurally similar to `fib` and `nqueens`, respectively, so we leave their detailed performance studies to Appendix A in the supplemental material.

The most obvious parameter affecting the performance is the size of the thread blocks used by our code transformations. Larger thread blocks clearly require more memory. More importantly, though, thread block size determines the fundamental tradeoff underlying our performance results. Larger block sizes lead to more work that can be vectorized, increasing SIMD utilization. However, large blocks suffer from poor locality, increasing cache misses. To achieve good performance, therefore, we want to achieve good SIMD utilization with the smallest possible block size.

SIMD utilization: Figure 8 shows how SIMD utilization changes with block size.[5] SIMD utilization is the percentage of tasks that are executed as part of full SIMD blocks—other tasks, which are part of the "epilog" of vectorized execution, lead to idle SIMD lanes. Higher SIMD utilization means more effective use of SIMD resources and, all else being equal, better performance. SIMD utilization for a benchmark is determined by vector width and block size, so for all benchmarks except `knapsack` and `uts`, utilization with respect to block size is the same for both platforms.

SIMD utilization increases rapidly with block size, and for all benchmarks, with or without re-expansion. Given a sufficiently large block, our transformations can achieve almost perfect utilization. Crucially, however, with re-expansion, the block size required for perfect utilization shrinks on

[5] In Figures 8–11, legends for `knapsack` apply to all graphs. "no reexp" refers to vectorization without re-expansion, while "reexp" includes our re-expansion technique.

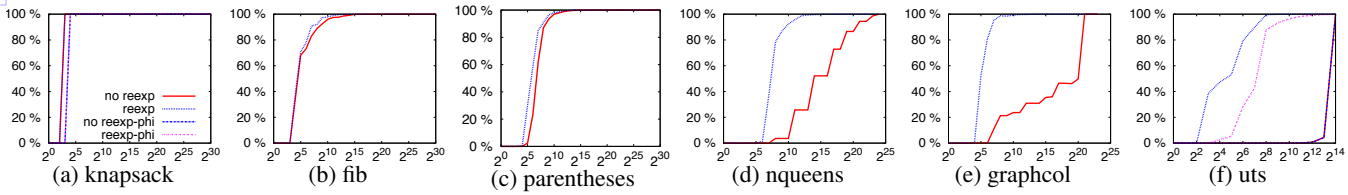

Figure 8: SIMD utilization. x-axis: block size; y-axis: percentage of tasks that can be vectorized

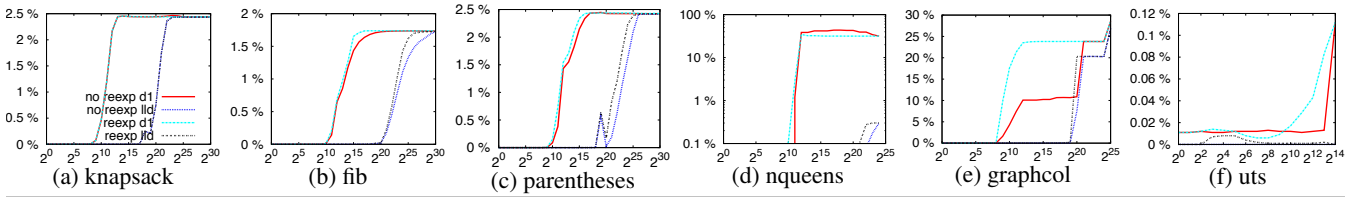

Figure 9: Xeon E5 cache miss rate. x-axis: block size; y-axis: miss rate for level 1 (d1) and last level (lld) caches

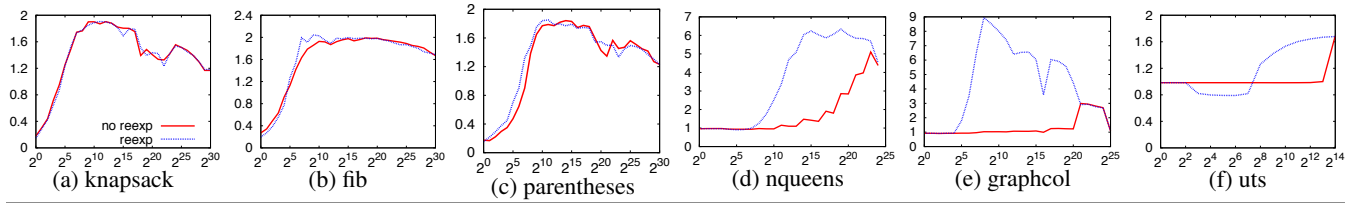

Figure 10: Xeon E5 speedup. x-axis: block size; y-axis: speedup relative to sequential baseline

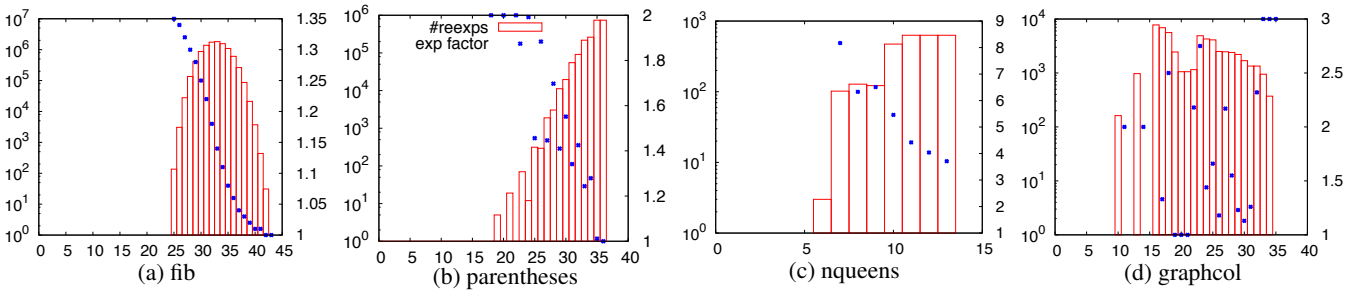

Figure 11: Benefits of re-expansion. x-axis: task level; left y-axis: number of re-expansions; right y-axis: factor of block size improvement due to re-expansion

several benchmarks (notably, `nqueens` and `graphcol`). To understand why, recall that without re-expansion, we generate parallel work using breadth-first expansion only at the beginning of the computation and the subsequent blocked depth-first execution cannot generate additional parallel work. Therefore, to achieve high utilization, we must generate a large amount of parallelism (large blocks) in the first depth first expansion before we begin depth-first execution. Re-expansion's ability to generate additional parallelism later in execution lets it tolerate a smaller block size. Re-expansion has little effect on utilization for some benchmarks, notably `knapsack`, `fib` and `parentheses`. For `knapsack`, re-expansion is never needed due to the perfectly balanced tree. The other two benchmarks have more subtle behavior, which we investigate more carefully later.

E5 cache efficiency and speedup: SIMD utilization only affects the amount of work that can be vectorized; it is not the only factor that affects performance. Another crucial factor, which militates against large blocks, is cache efficiency. It is the interplay between utilization and efficiency that determines speedup. We next investigate this behavior on the E5 platform. Appendix B in the supplemental material presents a similar investigation for the Xeon Phi.

Figure 9 shows both the L1 and last-level data cache misses rates with varying block size, with and without re-expansion. We see that as the block size grows, cache misses increase. To understand why, note that all of the threads in a thread block are accessed *twice*: once when they are added to the thread block, and a second time when they are executed. If the thread block is too large, the thread data will have been evicted by the second access. Unsurprisingly, we see fairly sharp discontinuities, representing cutoffs when blocks no longer fit in cache. Different benchmarks have fairly different cache behaviors, as they have different computational patterns; some benchmarks like `fib` do very little data access, while others like `nqueens` and `graphcol` perform lots of lookups. Nevertheless, the broad trend of increasing cache misses with block size persists.

Our vectorization speedup is due to a combination of both SIMD utilization and cache behavior. Figure 10 shows the

overall speedups of our techniques with varying block sizes. For all the benchmarks except `uts`, we see a consistent pattern: speedup increases with block size as SIMD utilization increases; then, at larger block sizes, cache misses begin to dominate, while we encounter diminishing utilization returns, causing speedups to drop.

These results demonstrate the key advantage of our re-expansion scheduling strategy. By generating more work throughout execution, re-expansion lets our transformed code achieve high SIMD utilization with smaller block sizes, letting us gain large benefits from vectorization before poor cache performance drags down overall speedups. This effect is most noticeable for `nqueens` and `graphcol`, where re-expansion achieves near-perfect SIMD utilization at block sizes small enough to avoid the cache-miss cliff, giving very high speedups. Even for benchmarks where re-expansion is not as critical, such as `fib` and `parentheses`, re-expansion lets us achieve peak speedup at somewhat smaller block sizes, reducing overall memory usage.

The exceptions to these trends are `knapsack` and `uts`. The former does not benefit from re-expansion due to its balanced computation tree; since threads never die out, the block size never gets small enough to trigger re-expansion. The latter has a relatively narrow computation tree and is very unbalanced, and hence performs best when the block size is large enough to obviate the need for ever doing depth-first execution in the first place (2^{14} threads).

Re-expansion benefit: We study the benefits from re-expansion in exposing data parallelism in Figure 11. For each level of the computation tree, the figure shows two quantities: the number of re-expansions performed at that level and the factor of increase in the number of tasks at the next level due to re-expansion. Larger factors denote greater benefit; a factor of 1 means that the block size did not change after re-expansion. We do not show `knapsack`'s and `uts` benchmarks since their execution never triggers re-expansion. Among the other benchmarks, re-expansion has limited benefit for `fib` and `parentheses` due to the fact that these computation trees are also relatively balanced and re-expansion is triggered fairly late, and does not generate much additional parallelism since the trees are not expanding (getting wider) any more. Re-expansion is much more useful in adapting to tree structures with base cases intermingled with recursive tasks at shallower depths. We observe this for `nqueens` and `graphcol`, which can get re-expansion factors as high as 8 and 3, respectively.

Benefits from stream compaction: We evaluate the benefits from stream compaction on two representative benchmarks: `fib`, one of the benchmarks with a small kernel, and `nqueens`, which has a larger kernel. Figure 12 shows the speedups achieved by the best block size configuration, as compared to the sequential execution, when the stream compaction is performed sequentially as compared to our table-lookup based compaction. We see that the table-lookup based compaction is faster in all cases, with significant improvements for smaller kernels. In fact, optimized stream compaction is crucial to performance on the smaller kernels. Even for benchmarks with larger kernels, we observe 5–10% overall performance improvement. We observe similar behavior for the other benchmarks considered as well.

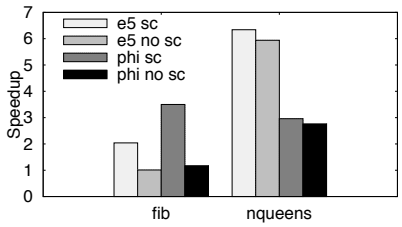

Figure 12: Speedup with and without stream compaction (sc) on E5 and Xeon Phi, normalized to sequential baseline

	Sequential		Vectorized		
Benchmark	Vect	non-Vect	Vect	non-Vect	Speedup
nqueens	0.94	0.06	0.06	0.03	10.74
graphcol	0.99	0.01	0.06	0.01	14.28
uts	0.81	0.19	0.20	0.20	2.50
minmax	0.62	0.38	0.04	0.25	3.48

Table 3: Estimated maximum vectorization speedup on E5

6.4 Opportunity analysis

Various factors preclude us from achieving the perfect speedup (i.e., 16 for 16-way SIMD) from vectorization. Here, we try to quantify the theoretically maximum achievable speedup. Given that only the kernel computation is vectorized, we compute the effect of Amdahl's law due to non-kernel overheads by looking at the number of kernel and non-kernel instructions. While number of instructions executed does not strictly determine performance, this opportunity study gives us some insight into vectorization potential (assuming 1.0 CPI). Since it is difficult to isolate the core computations in benchmarks with small tasks (`fib`, `parentheses`, `knapsack`, and `binomial`), we focus on the remaining benchmarks.

Table 3 shows the fraction of vectorizable and non-vectorizable instructions for the remaining benchmarks. The Sequential columns indicate that a significant fraction of computation is vectorizable. In our modeled vectorized code, we assume perfect speedup for the vectorizable instructions (column 4), reducing the instruction count by a factor of the vector width. We profile the re-expansion version of the code to account for changes in the number of non-vectorizable instructions due to our transformations (column 5). Note that our transformations can occasionally reduce the number of non-kernel instructions (e.g., `nqueens` and `minmax`) due to the way they optimize stack management operations. The modeled maximum speedup is the ratio of the total number of dynamic instructions in the modeled vectorized version to the sequential versions. We see that even with perfect vectorization, the anticipated speedup for `uts` and `minmax` is only 2.5 and 3.48 respectively (due to the

large number of non-kernel instructions which are not vectorized). `nqueens` and `graphcol` fare better. Comparing with Table 2, we see that our vectorized implementations achieve a large fraction of this theoretical max speedup despite suffering from overheads such as cache misses.

7. Related Work

Parallelism for multicores: Many modern programming languages for multicores, such as as Cilk family [Blumofe et al. 1995; Frigo et al. 1998; Danaher et al. 2006], Thread Building Blocks [Reinders 2007], Task Parallel Library [TPL], OpenMP [OpenMP Architecture Review Board 2008], X10 [X10], allow programmers to express task parallelism using constructs similar to our `spawn` directive. Two important variants of work-stealing schedulers are relevant to our work: As described in Section 2, the *work-first strategy* [Frigo et al. 1998] is similar to our depth-first strategy—when a processor spawns a task, it places the continuation on its local pool (to be potentially stolen by other processors) and immediately starts executing the newly spawned task. In contrast, the *help-first strategy* [Guo et al. 2009] is similar to breadth-first execution—a processor places the newly spawned task on its local pool and immediately executes the continuation. Guo et al. [2009] propose using help-first scheduling to generate work quickly, and work-first scheduling thereafter to bound space usage. This strategy is somewhat similar to the execution strategy adopted by our initial code transformations that begin with breadth-first execution then switch to depth-first execution, though a traditional work-stealing scheduler would not provide the necessary structured execution for vectorization.

Parallelism for vector units/GPUs: Modern vectorizing compilers attempt to automatically perform vectorization for small loops in programs using a variety of techniques [Nuzman and Zaks 2008; Maleki et al. 2011]. However, they tend to target programs written in a very structured, data-parallel manner, and cannot handle even moderately complex programs [Maleki et al. 2011]. In more restricted domains, there has been more success in SIMDizing programs through synthesis [Barthe et al. 2013] and code generation from domain-specific languages [Ren et al. 2013] and other restricted sets of problems [Kim and Han 2012; Jo et al. 2013]; these approaches do not work for more general programs. Most work in mapping complex applications to vector units has been done by hand [Chhugani et al. 2012; Dammertz et al. 2008; Havel and Herout 2010; Hernquist 1990; Kim et al. 2010].

GPUs offer a more programmable interface than vector units on CPUs, but the most common programming models for GPUs are fundamentally data-parallel [NVIDIA; Stone et al. 2010]. In recent years, several attempts have been made to take GPUs' inherently data-parallel execution model and adapt it to target task-parallel programs [Tzeng et al. 2010; Aila and Laine 2009; Huo et al. 2013]. Perhaps most related to our work, Orr et al. [2014] provide a hardware implementation of the *channels* model proposed by Gaster and Howes [2012] and provide a mapping from simple Cilk-style programs to their channels implementation. Interestingly, the execution model imposed by channels on these programs resembles the level-by-level breadth-first execution strategy of our initial code transformation. To control space, they propose another hardware modification that allows the execution of one level of computation to be suspended—in essence, only processing part of each level of the tree. An interesting avenue of future work would be to compare Orr et al.'s scheduling strategy with our proposed strategies.

The key distinctions between our work and this work on GPUs are: (1) GPUs provide hardware support for gather and scatter operations and execution masking, meaning that GPU approaches do not need to consider data and computation organization as carefully, and (2) the only GPU implementation that targets the similar fine-grained task parallelism as our techniques requires custom hardware support, and is not suitable for targeting commodity vector hardware.

Stream compaction for vectorization: Stream compaction was first introduced as a general technique for managing blocks of data operated on by vector operations by Ren et al. [2013]; they performed stream compaction for four-wide vector units, but did not describe a general approach for arbitrary-length vectors. Mytkowicz et al. [2014] described a general permutation strategy for block management. Permutation is a generalization of stream compaction. However, because stream compaction is a simpler problem, our algorithm is more efficient since it is linear in the stream size (rather than quadratic) and can trade-off between the size of pre-computed tables and the number of lookups.

8. Conclusions

Vectorizing task-parallel programs requires solving several critical challenges: finding data-parallelism for vectorization; controlling space usage; and ensuring that SIMD units stay fully utilized. We present code transformations and scheduling strategies that address these problems, allowing recursive, task-parallel programs to be mapped efficiently to commodity vector hardware. Moreover, our stream compaction algorithm is applicable beyond our block management code, and could be integrated in production compilers.

Our results represent a first cut at mapping task-parallel programs to processors with SIMD units, and there are many opportunities for improved performance. For example, the next version of the Xeon Phi will support character-level vector operations. With our general stream compaction implementation, our scheme will automatically be able to take advantage of the increased vector widths of new hardware. Moreover, while our current results focus on improving single-core performance by leveraging SIMD units, our programming model is a standard task-parallel language. As Appendix C shows, it is feasible to integrate multicore parallelism with traditional work-stealing and our SIMDization technology. We plan to investigate this hybrid further in future work.

References

T. Aila and S. Laine. Understanding the Efficiency of Ray Traversal on GPUs. In *Proceedings of the Conference on High Performance Graphics 2009*, HPG '09, pages 145–149, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-603-8. . URL http://doi.acm.org/10.1145/1572769.1572792.

Barcelona OpenMP Task Suite (BOTS). Barcelona OpenMP Task Suite (BOTS). https://pm.bsc.es/projects/bots.

G. Barthe, J. M. Crespo, S. Gulwani, C. Kunz, and M. Marron. From Relational Verification to SIMD Loop Synthesis. In *Proceedings of the 18th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP '13, pages 123–134, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-1922-5. . URL http://doi.acm.org/10.1145/2442516.2442529.

R. D. Blumofe, C. F. Joerg, B. C. Kuszmaul, C. E. Leiserson, K. H. Randall, and Y. Zhou. Cilk: An Efficient Multithreaded Runtime System. In *Proceedings of the Fifth ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPOPP '95, pages 207–216, New York, NY, USA, 1995. ACM. ISBN 0-89791-700-6. . URL http://doi.acm.org/10.1145/209936.209958.

J. Chhugani, C. Kim, H. Shukla, J. Park, P. Dubey, J. Shalf, and H. D. Simon. Billion-particle SIMD-friendly Two-point Correlation on Large-scale HPC Cluster Systems. In *SC*, 2012. ISBN 978-1-4673-0804-5. URL http://dl.acm.org/citation.cfm?id=2388996.2388998.

Cilk. Cilk. http://supertech.csail.mit.edu/cilk/.

H. Dammertz, J. Hanika, and A. Keller. Shallow Bounding Volume Hierarchies for Fast SIMD Ray Tracing of Incoherent Rays. In *Proceedings of the Nineteenth Eurographics Conference on Rendering*, EGSR '08, pages 1225–1233, Aire-la-Ville, Switzerland, Switzerland, 2008. Eurographics Association. . URL http://dx.doi.org/10.1111/j.1467-8659.2008.01261.x.

J. S. Danaher, I.-T. A. Lee, and C. E. Leiserson. Programming with Exceptions in JCilk. *Sci. Comput. Program.*, 63(2):147–171, Dec. 2006. ISSN 0167-6423. . URL http://dx.doi.org/10.1016/j.scico.2006.05.008.

J. O. Eklundh. A Fast Computer Method for Matrix Transposing. *IEEE Trans. Comput.*, 21(7):801–803, July 1972. ISSN 0018-9340. . URL http://dx.doi.org/10.1109/T-C.1972.223584.

M. Frigo, C. E. Leiserson, and K. H. Randall. The Implementation of the Cilk-5 Multithreaded Language. In *Proceedings of the ACM SIGPLAN 1998 Conference on Programming Language Design and Implementation*, PLDI '98, pages 212–223, New York, NY, USA, 1998. ACM. ISBN 0-89791-987-4. . URL http://doi.acm.org/10.1145/277650.277725.

M. Frigo, P. Halpern, C. E. Leiserson, and S. Lewin-Berlin. Reducers and Other Cilk++ Hyperobjects. In *Proceedings of the Twenty-first Annual Symposium on Parallelism in Algorithms and Architectures*, SPAA '09, pages 79–90, New York, NY, USA, 2009. ACM. ISBN 978-1-60558-606-9. . URL http://doi.acm.org/10.1145/1583991.1584017.

B. Gaster and L. Howes. Can GPGPU Programming Be Liberated from the Data-Parallel Bottleneck? *Computer*, 45(8):42–52, August 2012. ISSN 0018-9162. .

Y. Guo, R. Barik, R. Raman, and V. Sarkar. Work-first and Help-first Scheduling Policies for Async-finish Task Parallelism. In *Parallel Distributed Processing, 2009. IPDPS 2009. IEEE International Symposium on*, pages 1–12, May 2009. .

J. Havel and A. Herout. Yet Faster Ray-Triangle Intersection (Using SSE4). *IEEE Transactions on Visualization and Computer Graphics*, 16(3):434–438, May 2010. ISSN 1077-2626. . URL http://dx.doi.org/10.1109/TVCG.2009.73.

L. Hernquist. Vectorization of Tree Traversals. *J. Comput. Phys.*, 87(1):137–147, Mar. 1990. ISSN 0021-9991. . URL http://dx.doi.org/10.1016/0021-9991(90)90230-X.

X. Huo, S. Krishnamoorthy, and G. Agrawal. Efficient Scheduling of Recursive Control Flow on GPUs. In *Proceedings of the 27th international ACM conference on International conference on supercomputing*, ICS '13, pages 409–420, New York, NY, USA, 2013. ACM. ISBN 978-1-4503-2130-3. . URL http://doi.acm.org/10.1145/2464996.2479870.

Y. Jo and M. Kulkarni. Enhancing Locality for Recursive Traversals of Recursive Structures. In *Proceedings of the 2011 ACM international conference on Object oriented programming systems languages and applications*, OOPSLA '11, pages 463–482, New York, NY, USA, 2011. ACM. ISBN 978-1-4503-0940-0. . URL http://doi.acm.org/10.1145/2048066.2048104.

Y. Jo, M. Goldfarb, and M. Kulkarni. Automatic Vectorization of Tree Traversals. In *Proceedings of the 22nd international conference on Parallel architectures and compilation techniques*, PACT '13, pages 363–374, Piscataway, NJ, USA, 2013. IEEE Press. ISBN 978-1-4799-1021-2. URL http://dl.acm.org/citation.cfm?id=2523721.2523770.

C. Kim, J. Chhugani, N. Satish, E. Sedlar, A. D. Nguyen, T. Kaldewey, V. W. Lee, S. A. Brandt, and P. Dubey. FAST: Fast Architecture Sensitive Tree Search on Modern CPUs and GPUs. In *Proceedings of the 2010 ACM SIGMOD International Conference on Management of Data*, SIGMOD '10, pages 339–350, New York, NY, USA, 2010. ACM. ISBN 978-1-4503-0032-2. . URL http://doi.acm.org/10.1145/1807167.1807206.

S. Kim and H. Han. Efficient SIMD Code Generation for Irregular Kernels. In *Proceedings of the 17th ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming*, PPoPP '12, pages 55–64, New York, NY, USA, 2012. ACM. ISBN 978-1-4503-1160-1. . URL http://doi.acm.org/10.1145/2145816.2145824.

S. Maleki, Y. Gao, M. J. Garzarán, T. Wong, and D. A. Padua. An Evaluation of Vectorizing Compilers. In *Proceedings of the 2011 International Conference on Parallel Architectures and Compilation Techniques*, PACT '11, pages 372–382, Washington, DC, USA, 2011. IEEE Computer Society. ISBN 978-0-7695-4566-0. . URL http://dx.doi.org/10.1109/PACT.2011.68.

T. Mytkowicz, M. Musuvathi, and W. Schulte. Data-parallel Finite-state Machines. In *Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems*, ASPLOS '14, pages 529–542, New York, NY, USA, 2014. ACM. ISBN 978-1-4503-2305-5. . URL http://doi.acm.org/10.1145/2541940.2541988.

D. Nuzman and A. Zaks. Outer-loop Vectorization: Revisited for Short SIMD Architectures. In *Proceedings of the 17th Inter-*

national Conference on Parallel Architectures and Compilation Techniques, PACT '08, pages 2–11, New York, NY, USA, 2008. ACM. ISBN 978-1-60558-282-5. . URL http://doi.acm.org/10.1145/1454115.1454119.

NVIDIA. CUDA. http://www.nvidia.com/object/cuda_home_new.html. URL http://www.nvidia.com/object/cuda_home_new.html.

S. Olivier, J. Huan, J. Liu, J. Prins, J. Dinan, P. Sadayappan, and C.-W. Tseng. UTS: An Unbalanced Tree Search Benchmark. In *Proceedings of the 19th International Conference on Languages and Compilers for Parallel Computing*, LCPC'06, pages 235–250, Berlin, Heidelberg, 2007. Springer-Verlag. ISBN 978-3-540-72520-6. URL http://dl.acm.org/citation.cfm?id=1757112.1757137.

OpenMP Architecture Review Board. OpenMP Specification and Features. http://openmp.org/wp/, May 2008. URL http://openmp.org/wp/openmp-specifications.

M. S. Orr, B. M. Beckmann, S. K. Reinhardt, and D. A. Wood. Fine-grain Task Aggregation and Coordination on GPUs. In *Proceeding of the 41st Annual International Symposium on Computer Architecuture*, ISCA '14, pages 181–192, Piscataway, NJ, USA, 2014. IEEE Press. ISBN 978-1-4799-4394-4. URL http://dl.acm.org/citation.cfm?id=2665671.2665701.

M. Puschel, J. M. Moura, J. R. Johnson, D. Padua, M. M. Veloso, B. W. Singer, J. Xiong, F. Franchetti, A. Gacic, Y. Voronenko, et al. SPIRAL: Code Generation for DSP Transforms. *Proceedings of the IEEE*, 93(2):232–275, 2005.

J. Reinders. *Intel Threading Building Blocks: Outfitting C++ for Multi-Core Processor Parallelism*. O'Reilly, 2007.

B. Ren, T. Poutanen, T. Mytkowicz, W. Schulte, G. Agrawal, and J. R. Larus. SIMD Parallelization of Applications that Traverse Irregular Data Structures. In *Proceedings of the 2013 IEEE/ACM International Symposium on Code Generation and Optimization (CGO)*, CGO '13, pages 1–10, Washington, DC, USA, 2013. IEEE Computer Society. ISBN 978-1-4673-5524-7. . URL http://dx.doi.org/10.1109/CGO.2013.6494989.

M. Steffen and J. Zambreno. Improving SIMT Efficiency of Global Rendering Algorithms with Architectural Support for Dynamic Micro-Kernels. In *Proceedings of the 2010 43rd Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO '43, pages 237–248, Washington, DC, USA, 2010. IEEE Computer Society. ISBN 978-0-7695-4299-7. . URL http://dx.doi.org/10.1109/MICRO.2010.45.

J. E. Stone, D. Gohara, and G. Shi. OpenCL: A Parallel Programming Standard for Heterogeneous Computing Systems. *IEEE Des. Test*, 12(3):66–73, May 2010. ISSN 0740-7475. . URL http://dx.doi.org/10.1109/MCSE.2010.69.

TPL. The Task Parallel Library. http://msdn.microsoft.com/en-us/magazine/cc163340.aspx, Oct. 2007. URL http://msdn.microsoft.com/en-us/magazine/cc163340.aspx.

S. Tzeng, A. Patney, and J. D. Owens. Task Management for Irregular-parallel Workloads on the GPU. In *Proceedings of the Conference on High Performance Graphics*, HPG '10, pages 29–37, Aire-la-Ville, Switzerland, Switzerland, 2010. Eurographics Association. URL http://dl.acm.org/citation.cfm?id=1921479.1921485.

X10. The X10 Programming Language. www.research.ibm.com/x10/, Mar. 2006. URL www.research.ibm.com/x10/.

252

Subject expertise?	Passing knowledge	Expert
Relevant?	Uninteresting	Compelling
Sound?	Flawed	Sound
Accept?	No	Yes
Strength of conviction?	Ambivalent	Adamant

Points for

Points against

Questions for authors

Other notes