

How Should Compilers Represent Fork-Join Parallelism?

by

William S. Moses

Submitted to the Department of Electrical Engineering and Computer
Science

in partial fulfillment of the requirements for the degree of

Master of Engineering in Electrical Engineering and Computer Science

at the

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

June 2017

© Massachusetts Institute of Technology 2017. All rights reserved.

Author
Department of Electrical Engineering and Computer Science
May 25, 2017

Certified by
Charles E. Leiserson
Professor
Thesis Supervisor

Certified by
Tao B. Schardl
Postdoctoral Associate
Thesis Supervisor

Accepted by
Christopher Terman
Chairman, Masters of Engineering Thesis Committee

How Should Compilers Represent Fork-Join Parallelism?

by

William S. Moses

Submitted to the Department of Electrical Engineering and Computer Science
on May 25, 2017, in partial fulfillment of the
requirements for the degree of
Master of Engineering in Electrical Engineering and Computer Science

Abstract

This thesis explores how fork-join parallelism, as supported by concurrency platforms such as Cilk and OpenMP, can be embedded into a compiler's intermediate representation (IR). Mainstream compilers typically treat parallel linguistic constructs as syntactic sugar for function calls into a parallel runtime. These calls prevent the compiler from performing optimizations across parallel control constructs. Remedying this situation is generally thought to require an extensive reworking of compiler analyses and code transformations to handle parallel semantics.

Tapir is a compiler IR that represents logically parallel tasks asymmetrically in the program's control flow graph. Tapir allows the compiler to optimize across parallel control constructs with only minor changes to its existing analyses and code transformations. To prototype Tapir in the LLVM compiler, for example, the Tapir team added or modified about 6000 lines of LLVM's 4-million-line codebase. Tapir enables LLVM's existing compiler optimizations for serial code — including loop-invariant-code motion, common-subexpression elimination, and tail-recursion elimination — to work with parallel control constructs such as spawning and parallel loops. Tapir also supports parallel optimizations such as loop scheduling.

This research reported in this thesis represents joint work with Tao B. Schardl and Charles E. Leiserson.

Thesis Supervisor: Charles E. Leiserson

Title: Professor

Thesis Supervisor: Tao B. Schardl

Title: Postdoctoral Associate

Acknowledgments

This thesis is derived from a project done in collaboration with Tao B. Schardl and Charles E. Leiserson. Together, we form the “Tapir team,” which I will refer to in this document. TB and Charles were instrumental in helping talk through many of the ideas that inspired the following work. Much of this content was previously published at the 2017 Symposium on Principles and Practice of Parallel Programming (PPoPP)[55]. I would like to thank the PPoPP reviewers for their helpful comments and their selection as best paper.

I would also like to thank my professional colleagues. I would like to thank Tim Kaler and the students of the Fall 2016 MIT class 6.172/6.871 *Performance Evaluation of Software Systems* for their patience in using the Tapir/LLVM compiler throughout the semester and reporting bugs. Additional thanks to Shahin Kamali, Bradley Kuszmaul, Bojan Serafimov, Jiahao Li, Dougie Kogut, and the entire MIT Supertech research group for many helpful discussions. Further thanks to Larry Hardesty of the MIT News Office for asking questions that helping simplify Figure 1-3. I would also like to thank Johannes Doerfert, Simon Moll, Vikram Adve, and Hal Finkel for a number of discussions on expanding Tapir and trying to bring its ideas into LLVM.

Finally, I would like to thank my family without whose constant love and support I would never have been able to make it to where I am today. I would like my parents, John and Marina Moses, my sister Sophia Moses, my grandfather Panayoti Stefanidis, and my dog Patches.

This research was supported in part by NSF Grants 1314547 and 1533644, in part by a MIT CSAIL grant from Foxconn, in part by the Intelligence Advanced Research Projects Activity (IARPA) via Department of Interior/ Interior Business Center (DoI/IBC) contract number D16PC00002, and in part by an MIT EECS SuperUROP. The U.S. Government is authorized to reproduce and distribute reprints for Governmental purposes notwithstanding any copyright annotation thereon. Disclaimer: The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of IARPA, DoI/IBC, or the U.S. Government.

Contents

1	Introduction	9
1.1	Previous approaches	11
1.2	The Tapir approach	13
1.3	Ease of implementation	14
1.4	Expressiveness of Tapir	15
1.5	Serial semantics	16
1.6	Optimizations	16
1.7	Evaluation of Tapir/LLVM	16
1.8	Contributions	18
1.9	Outline	18
2	Tapir	19
2.1	Tapir instructions	19
2.2	Static single-assignment form	21
2.3	Asymmetry in Tapir	22
2.4	Parallel loops in Tapir	23
3	Analysis passes	24
3.1	Constraints on transformations	24
3.2	Alias analysis	25
3.3	Dominator analysis	25
3.4	Data-flow analysis	26

4	Optimization passes	28
4.1	Common-subexpression elimination	28
4.2	Loop-invariant code motion	29
4.3	Tail-recursion elimination	30
4.4	Parallel-loop scheduling and lowering	31
4.5	Other optimization passes	32
5	Auxiliary software	33
6	Evaluation	35
6.1	Benchmarking	37
6.2	Overall performance	38
7	Related work	41
8	Conclusion	43
A	Artifact description	45
A.1	Building Tapir/LLVM from source	46
A.2	Running the benchmark suite	47
A.3	Evaluation and expected result	48

List of Figures

1-1	A function that GCC, ICC, and Cilk Plus/LLVM all fail to optimize effectively.	10
1-2	Comparison between a traditional CFG with symmetric parallelism and Tapir's CFG with asymmetric parallelism.	12
1-3	Code changes required to implement the Tapir/LLVM prototype.	14
1-4	Graph of work efficiency comparison between Reference and Tapir/LLVM.	17
2-1	Tapir CFG for the parallel loops in Figure 1-1.	23
4-1	Example of common-subexpression elimination on a Cilk program.	29
4-2	Example of tail-recursion elimination on a parallel quicksort program.	30
6-1	The compilation pipelines for Clang/LLVM, Tapir/LLVM, and Reference.	36
6-2	Descriptions of the 20 benchmarks used to evaluate Tapir/LLVM	37
6-3	Numbers for performance comparison between Reference and Tapir/LLVM.	40

Chapter 1

Introduction

Mainstream compilers, such as GCC [62], ICC [18], and LLVM [26] provide linguistic extensions for frameworks such as Cilk Plus[16] and OpenMP [4, 47] that allow programmers to write fork-join parallel programs. Typically in such frameworks, one can specify parallelism at a high level by denoting tasks or loops iterations that may be executed concurrently.

Although these mainstream compilers support fork-join parallelism, they struggle to optimize programs when they encounter such linguistic constructs. Paradoxically this can even mean that programs you'd expect to show large parallel speedups, are slower than the equivalent serial code. Consider, for example, the parallel `cilk_for` loop on lines 6–7 in Figure 1-1a, which indicates that iterations of the loop are free to execute in parallel. In a serial version of this loop, where the `cilk_for` keyword is replaced by an ordinary `for` keyword, each of the compilers GCC 5.3.0, ICC 16.0.3, and Cilk Plus/LLVM 3.9.0 observes that the call to `norm` on line 7 produces the same value in every iteration of the loop, and they optimize the loop by computing this value only once before the loop executes. This optimization dramatically reduces the total time to execute `normalize` from $\Theta(n^2)$ to $\Theta(n)$. Although this same optimization can, in principle, be performed on the actual parallel loop in the figure, no mainstream compiler performs this code-motion optimization. The same is true when the parallel loop is written using OpenMP, as shown in Figure 1-1b.

This failure to optimize stems from how these compilers for serial languages implement parallel linguistic constructs. The compiler for a serial language, such as C [22] or

a

```

01 __attribute__((const))
02 double norm(const double *A, int n);
03
04 void normalize(double *restrict out,
05               const double *restrict in, int n) {
06     cilk_for (int i = 0; i < n; ++i)
07         out[i] = in[i] / norm(in, n);
08 }

```

b

```

09 __attribute__((const))
10 double norm(const double *A, int n);
11
12 void normalize(double *restrict out,
13               const double *restrict in, int n) {
14     #pragma omp parallel for
15     for (int i = 0; i < n; ++i)
16         out[i] = in[i] / norm(in, n);
17 }

```

Figure 1-1: A function that GCC, ICC, and Cilk Plus/LLVM all fail to optimize effectively. **a** A Cilk version of the code. The `cilk_for` loop on lines 6–7 allows each iteration of the loop to execute in parallel. The `norm` function computes the norm of a vector in $\Theta(n)$ time. The call to `norm` on line 7 can be safely moved outside of the loop, but none of these three mainstream compilers perform this code motion, even though they all do so when the `cilk_for` keyword is replaced with an ordinary `for` keyword. **b** The corresponding OpenMP code.

C++ [63], can be viewed as consisting of three phases: a front end, a middle end, and a back end. The front end parses and type-checks the input program and translates it to an *intermediate representation (IR)*, which represents the control flow of the program as a more-or-less language-independent *control-flow graph (CFG)* [2, Sec. 8.4.3]. The middle end consists of optimization passes that transform the IR into a more-efficient form. These optimizations tend to be independent of the instruction-set architecture of the target computer. The back end translates the optimized IR into machine code, performing low-level machine-dependent optimizations.

GCC, ICC, and Cilk Plus/LLVM all *lower* the parallel constructs — transform the parallel constructs to a more-primitive representation — in the front end. To compile the code in Figure 1-1a, for example, the front-end translates the parallel loop in lines 6–7 into IR in two steps. (The OpenMP code in Figure 1-1b is handled similarly.) First, the loop body (line 7) is lifted into a helper function. Next, the loop itself is replaced with a call to a library function implemented by the Cilk Plus runtime system, which takes as arguments the loop bounds and helper function, and handles the spawning of the loop iterations for parallel execution. Since this process occurs in the front end, it renders the parallel loop unrecognizable to middle-end loop-optimization passes, such as code motion. In short, these compilers treat parallel constructs as syntactic sugar for opaque runtime calls, which confounds the many middle-end analyses and optimizations.

1.1 Previous approaches

This thesis aims to enable middle-end optimizations involving fork-join control flow by embedding parallelism directly into the compiler IR, an endeavor that has historically been challenging [31, 30]. For example, it is well documented [40] that traditional compiler transformations for serial programs can jeopardize the correctness of parallel programs. In general, four types of approaches have been proposed to embed parallelism in a mainstream compiler IR.

First, the compiler can use metadata to delineate logical parallelism. LLVM’s parallel loop metadata [34], for example, is attached to memory accesses in a loop to indicate that they have no dependence on other iterations of the same loop. LLVM can only conclude that a loop is parallel if all its memory accesses are labeled with this metadata. Unfortunately, encoding parallel loops in this way is fragile, since a compiler transformation that moves code into a parallel loop risks serializing the loop from LLVM’s perspective.

Second, the compiler can use intrinsic functions to demark parallel tasks. (For examples, see [65, 48, 32].) Often, either existing serial analyses and optimizations must be shut down when code contains these intrinsics, or the intrinsics offer minimal opportunities for compiler optimization.

Third, the compiler can use a separate IR to encode logical parallelism in the program. The HPIR [65, 6], SPIRE [23], and INSPIRE [21] representations, for instance, model parallel constructs using an alternative IR, such as one based on the program’s abstract syntax tree [2, Sec. 2.5.1]. Such an IR can support optimizations involving parallel constructs without requiring changes to existing middle-end optimizations. But adopting a separate IR into a mainstream compiler has historically been criticized [33] as requiring considerable effort to engineer, develop, and maintain the additional IR to the same standards as the compiler’s existing serial IR.

Fourth, the compiler can augment its existing IR to encode logical parallelism, which is the approach that Tapir follows. Unlike Tapir, all prior research on parallel precedence graphs [61, 60], parallel flow graphs [59, 13], concurrent control-flow graphs [28, 46], and parallel program graphs [53, 52] represent parallel tasks as symmetric entities in a CFG.

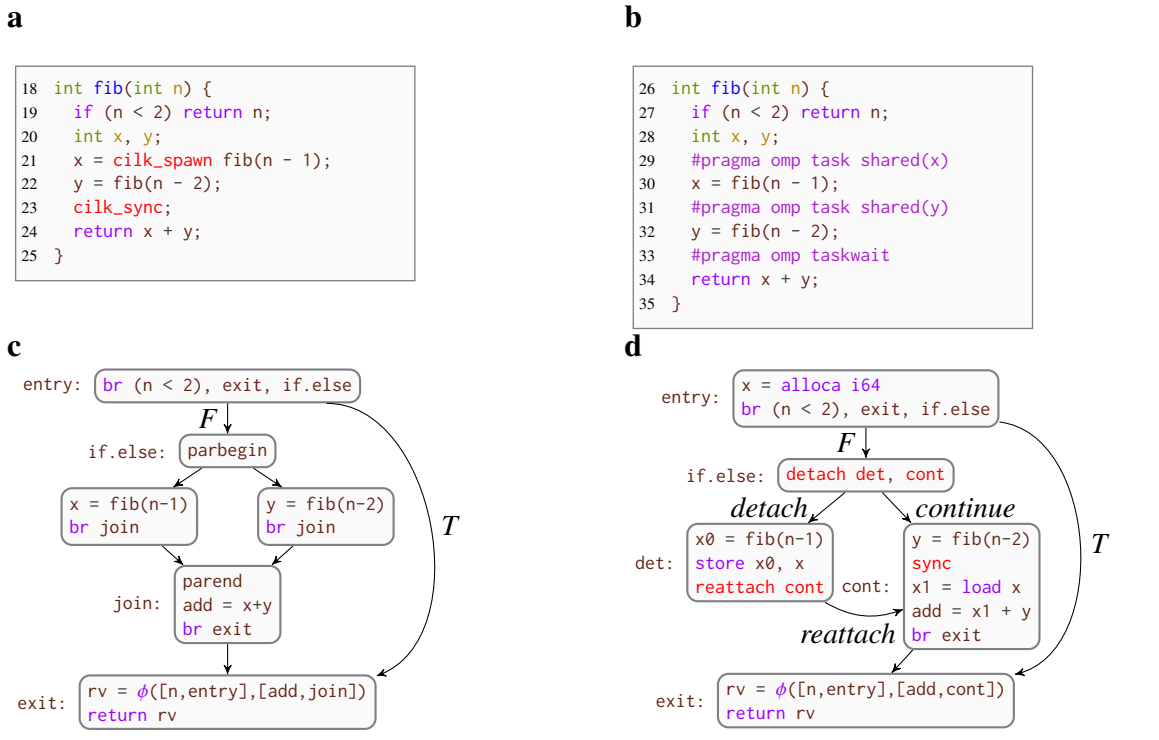


Figure 1-2: Comparison between a traditional CFG with symmetric parallelism and Tapir’s CFG with asymmetric parallelism. **a** The Cilk function `fib` computes Fibonacci numbers. The `cilk_spawn` on line 21 allows the two recursive calls to `fib` to execute in parallel, and the `cilk_sync` on line 23 waits for the spawned call to return. A serial execution of `fib` executes `fib(n-1)` before `fib(n-2)`. **b** A comparable implementation of `fib` using OpenMP task parallelism. **c** A CFG for `fib` that encodes parallelism symmetrically. Rectangles denote basic blocks, which contain C-like pseudocode for `fib`. Edges denote control flow between basic blocks. The `parbegin` and `parend` statements create and synchronize the parallel calls to `fib`. The `br` instruction encodes either an unconditional or a conditional branch. True and false edges from a conditional branch are labeled *T* and *F*, respectively. The ϕ instruction, used to support a static-single-assignment (SSA) form of the program (see Chapter 2), takes as its arguments pairs that associate a value with each predecessor basic block of the current block. At runtime the ϕ instruction returns the value associated with the predecessor basic block that executed immediately before the current block. **d** The Tapir CFG for `fib`, which encodes parallelism asymmetrically. The `alloca` instruction allocates shared-memory storage on the call stack for a local variable. Chapter 2 defines the `detach`, `reattach`, and `sync` instructions and the `detach`, `reattach`, and `continue` edge types.

For the parallel `fib` function in Figures 1-2a and 1-2b, for example, the parallel flow graph in Figure 1-2c illustrates how forked subcomputations might be represented symmetrically. Some of these approaches struggle to represent common parallel constructs, such as parallel loops [28, 23], while others exhibit problems when subjected to standard compiler

analyses and transformations for serial programs [28, 52, 13, 25, 61, 60, 51]. Existing serial-program analyses in LLVM, for example, assume that a basic block with multiple predecessors can observe the variables of only one predecessor at runtime. For the parallel flow graph in Figure 1-2c, however, instructions in the `join` block must observe the values of `x` and `y` from both of its predecessors, as has been observed by [28]. Parallel loops exacerbate this problem by allowing a dynamic number of tasks to join at the same basic block. Previous research [51, 1] has proposed solutions to these problems, including additional representations of the program and augmented analyses that account for interleavings of parallel instructions, but adopting these techniques into a mainstream compiler seems to require extensive changes to the existing codebase.

1.2 The Tapir approach

This thesis introduces Tapir, a compiler IR that represents logical fork-join parallelism asymmetrically in the program’s CFG. The asymmetry corresponds to the assumption of *serial semantics* [12], which means it is always semantically correct to execute parallel tasks in the same order as an ordinary serial execution.

Tapir adds three instructions — `detach`, `reattach`, and `sync` — to the IR of an ordinary serial compiler to express fork-join parallel programs with serial semantics. Figure 1-2d illustrates the Tapir CFG for the `fib` function. As with the symmetric parallel flow graph in Figure 1-2c, Tapir places the logically parallel recursive calls to `fib` in separate basic blocks. But these blocks do not join at a synchronization point symmetrically. Instead, one block connects to the other, reflecting the serial execution order of the program.

The Tapir approach provides five advantages:

1. Introducing fork-join parallelism into the compiler is relatively easy.
2. The IR is expressive and can represent fork-join control constructs from different parallel-language extensions.
3. Tapir parallel constructs harmonize with the invariants associated with existing representations of serial code.
4. Standard serial optimizations work on parallel code with few modifications.

- The optimizations enabled by Tapir’s parallelism constructs are effective in practice.

I discuss each of these advantages in turn.

1.3 Ease of implementation

Tapir’s asymmetric representation of logically parallel tasks makes it relatively simple to integrate Tapir into an existing compiler’s intermediate representation such as LLVM IR [34]. Figure 1-3 documents the lines of code added, modified, or deleted to implement a prototype of Tapir in LLVM. As Figure 1-3 shows, Tapir/LLVM was implemented with about 6000 lines, compared to LLVM’s roughly 4-million-line codebase. Moreover, fewer than 2000 lines of code were needed to adapt LLVM’s existing compiler analyses and transformations to accommodate Tapir.

<i>Compiler Component</i>	<i>LLVM 4.0svn</i>	<i>Tapir/LLVM</i>	
Instructions	105,995	943	} 1,768
Memory Behavior	21,788	445	
Optimizations	152,229	380	
Parallelism Lowering	0	3,782	
Other	3,803,831	460	
Total	4,083,843	6,010	

Figure 1-3: Breakdown of the lines of code added, modified, or deleted in LLVM to implement the Tapir/LLVM prototype.

The breakdown of lines is as follows. The lines for “Instructions” add Tapir’s instructions to LLVM IR and adapt LLVM’s routines for reading and writing LLVM IR and bit-code files. Conceptually, these changes allow LLVM to correctly compile a Tapir program to a serial executable with no optimizations. The lines for “Memory Behavior” control how Tapir instructions may interact with memory operations, preventing the compiler from creating any races. The lines for “Optimizations” perform any adjustments required for LLVM analyses and transformations to compile a Tapir program at optimization level `-O3`. Most of these modifications are not necessary for creating a correct executable but are added to allow the compiler to perform additional optimizations, such as parallel tail-recursion elimination (described in Chapter 4). The lines for “Parallelism Lowering” translate Tapir

instructions into Cilk Plus runtime calls and allow the code to be race-detected with a provably good race detector [11]. The lines for “Other” address a bug in LLVM’s implementation of `setjmp` and implement useful features for our development environment.

1.4 Expressiveness of Tapir

Tapir can express logical fork-join parallelism in parallel programs that have serial semantics. For example, Figure 1-2 illustrates how Tapir can express the parallelism encoded by the `cilk_spawn` and `cilk_sync` linguistics from Cilk++ [29] and Cilk Plus [16], as well as the parallelism encoded by OpenMP `task` and `taskwait` clauses [4]. Similarly, Tapir can express the parallelism encoded by OpenMP parallel sections [47] and Habanero’s `async` and `finish` constructs [8]. Tapir can also express parallel loops, including `cilk_for` loops and OpenMP parallel loops that have serial semantics (described in Chapter 2). Other parallel constructs, such as those proposed in the C++17 parallelism extensions, can be represented as well. However, parallel operations that cannot be expressed in terms of fork-join parallelism, such as OpenMP’s `ordered` clause, cannot be represented directly using Tapir’s `detach`, `reattach`, and `sync` instructions.

Tapir makes minimal assumptions about the consistency [49, 7] of concurrent memory accesses. Tapir assumes that memory is shared among parallel tasks and that virtual-register state is local to each task. Parallel instructions in Tapir can exhibit a *determinacy race*¹ [11] if they access the same memory location concurrently and at least one instruction writes to that location. Tapir itself does not fully define the possible outcomes of a determinacy race, and instead defers to existing compiler mechanisms, such as LLVM’s atomic memory-ordering constraints [34], to define whichever memory model they choose. For any targeted runtime system, Tapir relies on a correct implementation of lowering in order to implement the necessary synchronization, but Tapir is oblivious to how that runtime system implements the synchronization.

¹Determinacy races are also called general races [43] and are distinct from data races, which involve nonatomic accesses to critical regions.

1.5 Serial semantics

By grounding its model of parallelism in serial semantics, Tapir enables common compiler optimizations for serial code to work on parallel code. Intuitively, because Tapir always allows parallel tasks to execute in their ordinary serial execution order, the compiler can to optimize parallel code in any manner that preserves the serial semantics of the program and does not introduce new determinacy races. These mild constraints support common optimizations on parallel code, such as sequentialization, which can be invalid under models of parallelism without serial semantics [64].

1.6 Optimizations

In practice, the Tapir team has found that Tapir enables a wide variety of standard compiler optimizations to work with parallel code. The prototype implementation of Tapir/LLVM, for example, successfully moves the call to `norm` in Figure 1-1 outside of the loop, just as it would for a serial `for` loop. As Chapter 4 discusses, Tapir enables other optimizations, including common-subexpression elimination [41, Sec. 12.2], loop-invariant-code motion [41, Sec. 13.2], and tail-recursion elimination [41, Sec. 15.1], to work on parallel code. Tapir also enables new optimizations on parallel control flow.

1.7 Evaluation of Tapir/LLVM

The compiler optimizations that Tapir enables are effective in practice. We evaluated the Tapir approach by measuring the performance of 20 Cilk application benchmarks compiled using Tapir/LLVM. We compared the performance of these executables to those produced by a comparable reference compiler, called Reference. Conceptually, Reference lowers parallel linguistic constructs directly into runtime calls, as mainstream compilers do today, but otherwise performs the same set of optimization passes as Tapir/LLVM. Chapter 6 describes our experimental setup in detail, including the design of Reference.

Figure 1-4 presents the results of comparing Tapir/LLVM and Reference in terms of the




Figure 1-4: Comparison of the work efficiency of 20 parallel application benchmarks compiled using Tapir/LLVM (X’s) and the comparable Reference compiler (O’s), described in Chapter 6, which lowers parallelism in the compiler front end. Each point plots the work efficiency T_S/T_1 of a compiled benchmark, where T_1 is the work of the benchmark and T_S is the running time of the serial elision of the benchmark. Higher values indicate better work efficiency. The horizontal line at 1.0 plots the theoretically maximum work efficiency $T_S/T_1 = 1$. Benchmarks are sorted by decreasing difference in work efficiency between Tapir/LLVM- and Reference-compiled executables. Benchmarks marked with an “L” use parallel loops, and benchmarks marked with an “S” use `cilk_spawn`.

“work efficiency” of the compiled benchmarks. To perform this comparison, We compiled each benchmark using each compiler and then ran the executable on a single processing core of a multicore machine to measure its *work*, the 1-core running time, denoted T_1 . We also used each compiler to compile, run, and measure the 1-core running time of the *serial elision* [12] of each benchmark, denoted T_S , in which the benchmark is converted into a corresponding serial program by replacing all parallel linguistic constructs with their serial equivalents. We then computed the *work efficiency* of each compiled benchmark, which is the ratio T_S/T_1 of the running time T_S of the benchmark’s serial elision divided by the work T_1 of the benchmark. In theory, the maximum possible work efficiency is $T_S/T_1 = 1$, but in practice, quirky behaviors of the compiler and multicore architecture can occasionally produce work efficiencies greater than 1. As Figure 1-4 shows, for most benchmarks, the executables compiled using Tapir/LLVM achieve equal or higher work efficiency than those compiled using Reference. Moreover, for many benchmarks, and particularly those implemented using parallel loops, Tapir/LLVM produces executables that achieve nearly optimal work efficiency. Chapter 6 elaborates on these experiments.

1.8 Contributions

This thesis makes the following research contributions:

- The design of a compiler IR that represents fork-join parallelism asymmetrically, which enables existing serial optimizations to operate on parallel code and which also enables parallel optimizations.
- The implementation of Tapir/LLVM in the LLVM compiler by modifying about 6000 source lines of code (0.15% of the 4-million-line LLVM codebase).
- The implementation of parallel optimizations such as unnecessary synchronization elimination and parallel-loop scheduling.
- Experiments that demonstrate the advantage of embedding fork-join parallelism into a compiler’s IR, as opposed to dealing with parallelism only in the compiler’s front end.

1.9 Outline

The remainder of this thesis is organized as follows. Chapter 2 describes Tapir’s representation and properties. Chapter 3 discusses how analysis passes can be adapted to operate on Tapir programs. Chapter 4 describes various optimizations on parallel control flow that Tapir enables. Chapter 5 describes auxiliary software we developed to exercise and test Tapir/LLVM. Chapter 6 discusses our evaluation of the effectiveness of Tapir. Chapter 7 discusses related work. Chapter 8 provides some concluding remarks. An appendix describes how to set up Tapir/LLVM and how to download and run our suite of application benchmarks.

Chapter 2

Tapir

This chapter describes how Tapir represents logically parallel tasks asymmetrically in the CFG of a program. I define Tapir’s three new instructions and how they interact with LLVM’s static single-assignment (SSA) form [2, Sec. 6.2.4]. Although I describe Tapir as an extension to LLVM IR [34], the Tapir team sees no reason why other compilers cannot gain similar advantages from Tapir-like instructions.

Like LLVM IR, Tapir treats a program function as a CFG $G = (V, E, v_0)$, where

- the set V of vertices represents the function’s *basic blocks*: sequences of LLVM instructions, where control flow can only enter through the first instruction and leave from the last instruction;
- the set E of edges denote control flow between (basic) blocks; and
- the designated vertex $v_0 \in V$ represents the *entry point* of the function.

2.1 Tapir instructions

Tapir extends LLVM IR with three instructions: `detach`, `reattach`, and `sync`. The `detach` and `reattach` instructions together delineate logically parallel tasks, and the `sync` instruction imposes synchronization on parallel tasks. The three instructions have the following syntax, where $b, c \in V$:

```
detach label b, label c
reattach label c
sync
```

The `label` keywords indicate that b and c are (labels of) basic blocks in V .

The `detach` and `reattach` instructions together delineate a parallel task as follows. A `detach` instruction terminates the block a that contains it and takes a *detached* block b and a *continuation* block c as its arguments. The `detach` instruction *spawns* the task starting at block b , allowing that task to execute in parallel with block c . The control-flow edge $(a, b) \in E$ is a *detach* edge, and the edge $(a, c) \in E$ is a *continue* edge. A `reattach` instruction, meanwhile, terminates the block a' that contains it and takes a single *continuation* block c as its argument, inducing a *reattach* edge $(a', c) \in E$ in the CFG. The `reattach` terminates the task spawned by a preceding `detach` instruction with the same continuation block. Together, a `detach` instruction and associated `reattach` instructions demark the start and end of a parallel task and indicate that that task can execute in parallel with their common continuation block.

For the example in Figure 1-2d, the `detach` in the `if.else` block and the `reattach` in the `det` block share the same continuation block `cont`. Together, this `detach` and this `reattach` indicate that the `det` block is a parallel task which can execute in parallel with the `cont` block. In general, a parallel task delineated by `detach` and `reattach` can consist of many basic blocks in a single-entry subgraph.

The `detach` and `reattach` instructions in a CFG obey several structural properties. A `reattach` instruction j *reattaches* a `detach` instruction i if i and j share a common continuation block and there is a path from the detached block of i to j . Tapir assumes that every CFG $G = (V, E, v_0)$ obeys the following invariants on every `detach` instruction i and `reattach` instruction j in G :

1. A `reattach` instruction reattaches exactly one `detach` instruction.
2. If j reattaches i , then every path from v_0 to the block terminated by j passes through the detach edge of i , that is, the detach edge of i *dominates* j .
3. Every path starting from the detached block of i must reach a block terminated by a `reattach` instruction that reattaches i .

4. If j reattaches i and a path from i to j passes through the detach edge of another `detach` instruction i' , then it must also pass through a `reattach` instruction j' that reattaches i' .
5. Every cycle containing a `detach` instruction i must pass through a `reattach` instruction that reattaches i .
6. The continuation block of j cannot contain any ϕ instructions [2, Sec 6.2.4].

These invariants imply that, at runtime, a `detach` instruction i with detached block b and continuation block c spawns the execution of a *detached sub-CFG*, which is the single entry sub-CFG starting at b induced by all blocks on paths from b to a `reattach` instruction that reattaches i .

The dynamic execution of the program organizes memory as a tree of *parallel contexts*. A new parallel context is created as a child of the current context when control enters a function or follows a detach edge. When control executes a `reattach` instruction or leaves a function, the context is destroyed and the parent's context becomes the current context. An `alloca` instruction allocates shared memory in the current context.

The `sync` instruction synchronizes tasks spawned within its parallel context. At runtime, a `sync` instruction dynamically waits for the set of sub-CFG's detached in the same parallel context or any of its descendant parallel contexts to reach a `reattach` instruction. In the Tapir CFG illustrated in Figure 1-2d, for example, the `sync` instruction in the `cont` block simply waits for the execution of the `det` block to complete. Unlike `reattach` instructions, `sync` instructions are not explicitly associated with `detach` instructions, and they, in fact, can be executed within conditionals. A `sync` instruction j *syncs* a `detach` instruction i if i and j belong to the same parallel context and the CFG detached by i cannot be guaranteed to have completed when j executes.

2.2 Static single-assignment form

LLVM's static single-assignment (SSA) form [2, Sec. 6.2.4] must be adapted for Tapir programs. SSA form ensures that each virtual register is set at most once in a function. LLVM IR employs the ϕ *instruction* [2, Sec 6.2.4] to combine definitions of a variable

from different predecessors of a basic block. In adapting SSA to Tapir, one concern is that a ϕ instruction might allow registers defined in the detached sub-CFG to be used in the continuation. A basic block containing a ϕ instruction must avoid inheriting register definitions from predecessors that are connected by reattach edges. Otherwise, a register in the detached sub-CFG might not have been computed by the time the continuation executes.

We implemented this constraint by simply forbidding reattach edges from going into basic blocks with ϕ instructions. But what if the continuation c of a `detach` instruction begins with a ϕ instruction? In this case, Tapir creates a new basic block c' containing only a branch instruction to c . Tapir reroutes the reattach and continuation edges originally going to c so that they go instead to c' . All other edges going to c are left in place.

The reason this solution works is as follows. No reattach edges in the resulting CFG go to blocks containing ϕ instructions. Because a detached sub-CFG does not dominate any outside block, registers in the detached CFG can only be used in ϕ instructions of the immediate successors of the detached sub-CFG. Since the continuation is the only immediate successor of the detached sub-CFG and it contains no ϕ instructions, no registers from the detached sub-CFG may be accessed in the continuation.

2.3 Asymmetry in Tapir

The `detach` and `reattach` instructions express parallel tasks asymmetrically both syntactically in the structure of the CFG and semantically in the way memory state is managed. Both asymmetries are illustrated in Figure 1-2d.

First, the CFG detached by a `detach` instruction is connected by a reattach edge to the continuation block of that instruction, even though they can execute in parallel. For example, the reattach edge between `det` and `cont` in Figure 1-2d breaks the symmetry between them. Reattach edges reflect the serial semantics of a Tapir program, which dictates that a serial execution of the program executes the detached CFG to completion before starting to execute the continuation block. In fact, the parallel task delineated by a `detach` and a `reattach` instruction can be *serialized* by replacing the `detach` instruction with an unconditional branch to its detached block and replacing the `reattach` with an unconditional

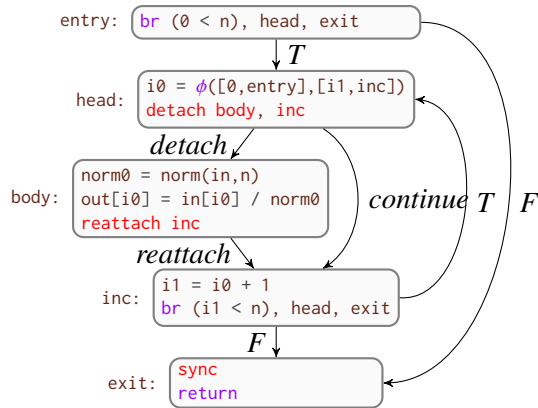


Figure 2-1: Tapir CFG for the parallel loops in Figure 1-1.

branch to its continuation block. In contrast, parallel flow graphs and similar previously explored representations join logically parallel tasks in the CFG at a synchronization point. By supporting separate `reattach` and `sync` instructions, Tapir decouples the termination of a parallel task from its synchronization.

Second, although memory state is shared among all parallel tasks in Tapir, a virtual register defined in a detached sub-CFG is not accessible in its parent parallel context. For example, the continuation block `cont` in Figure 1-2d cannot assume that the register value `x0` returned by `fib(n-1)` in block `det` is accessible, because the two basic blocks belong to different parallel contexts. Thus, `cont` must load it again after the `sync` instruction.

2.4 Parallel loops in Tapir

Figure 2-1 illustrates Tapir’s default representation of the parallel loops from Figure 1-1. As Figure 2-1 shows, Tapir can represent a parallel loop in the CFG as an ordinary loop, where the `head` block repeatedly spawns the `body` block, and the `exit` block syncs the detached CFG’s. Chapter 4 describes how this representation of parallel loops allows existing compiler loop optimizations to operate on Tapir parallel loops with only minor modifications. Although this loop structure can exhibit poor parallel performance when the loop body is small, separate optimization passes in Tapir/LLVM (see Chapter 4) transform this parallel-loop representation into a divide-and-conquer form that exhibits good performance.

Chapter 3

Analysis passes

This chapter describes how LLVM’s analysis passes can be adapted to operate on Tapir programs. I first discuss constraints on how Tapir programs can be safely transformed. Implementing these constraints on LLVM optimization passes primarily involves adapting standard compiler analyses — specifically alias analysis [2, Ch. 12], dominator analysis [2, Ch. 9], and data-flow analysis [2, Ch. 9] — to accommodate Tapir’s instructions. I describe how each of these analyses was minimally modified to support Tapir.

3.1 Constraints on transformations

To be correct, a code transformation on a Tapir program must preserve the program’s serial semantics, and it must not introduce any new behaviors into the program’s set of behaviors. A program can exhibit more than one behavior if it contains a determinacy race. In general, the result of a determinacy race can vary nondeterministically from run to run depending on the order in which the participating instructions access the memory location. To avoid introducing new behaviors, code transformations must not create determinacy races, although they can eliminate determinacy races. Many existing serial optimizations can be adapted to respect these properties by adapting the standard compiler analyses they rely on. I now describe how LLVM’s alias, dominator, and data-flow analyses were adapted for Tapir.

3.2 Alias analysis

LLVM uses alias analysis [2, Ch. 12] to determine whether different instructions might reference the same locations in memory, and in particular, to restrict the reordering of instructions that access the same memory. Tapir/LLVM modifies LLVM's alias analysis to prevent optimizations that move code around from introducing determinacy races. In particular, Tapir adapts LLVM's alias analysis to treat the instructions as if they access memory. For example, consider an instruction k that performs a load or a store. There are four cases to consider when moving k around either a `detach` instruction i or a `sync` instruction j :

1. The instruction k moves from before i to after i .
2. The instruction k moves from after i to before i .
3. The instruction k moves from before j to after j .
4. The instruction k moves from after j to before j .

Neither Case 2 nor Case 3 can introduce a determinacy race, because both motions serialize the execution of k with respect to the sub-CFG detached by i . Cases 1 and 4 might introduce a determinacy race, however, if k loads or stores a memory location that is also accessed by the CFG detached by i . To handle Case 1, i is treated as if it were a function call that accesses all memory locations accessed in the CFG detached by i . Similarly, for Case 4, j is treated as if it were a function call that accesses all memory locations accessed by all instructions that j might sync. A `reattach` instruction is treated as a compiler fence that prevents instructions from moving across it. With these modifications, existing rules in LLVM that restrict reordering of loads and stores properly restrict memory reordering around Tapir's instructions.

3.3 Dominator analysis

Optimization passes determine what values are available to an instruction in part by using dominator analysis [2, Ch. 9], which deduces the dominance relation between all basic blocks and edges in a CFG. To handle Tapir programs correctly, optimization passes must

not mistakenly cause instructions to use virtual registers that are defined in logically parallel tasks. If instruction i dominates instruction j , then an optimization pass might assume that the value produced by i is always available when j executes.

The asymmetry of Tapir’s representation allows LLVM’s dominator analysis to analyze Tapir programs correctly without any changes. Ignoring the names of edges, the difference between the CFG $G = (V, E, v_0)$ of a Tapir program and the CFG $G' = (V, E', v_0)$ of its serial elision is the set $E - E'$ of continue edges, each of which connects a `detach` instruction to its continuation. A continue edge short-cuts a detached sub-CFG, changing the continuation’s immediate dominator from the detached sub-CFG to the block containing the `detach` instruction itself. This configuration of `detach`, `reattach`, and continue edges looks much like an ordinary `if` construct in which the detached sub-CFG is conditionally executed. As a result, dominator analysis never concludes that an instruction in a detached sub-CFG can execute before the corresponding continuation block.

3.4 Data-flow analysis

A wide class of code transformations, including those that might move instructions across a `reattach` edge, rely on data-flow analysis [2, Ch. 9] to examine the propagation of values along different paths through a CFG $G = (V, E, v_0)$. Fundamental to data-flow analysis is an understanding of the set of possible program states at the beginning and end of each basic block $b \in V$, denoted $\text{IN}(b)$ and $\text{OUT}(b)$, respectively.

To illustrate how LLVM’s data-flow analyses were adapted to Tapir, let us examine the particular case of forward data-flow analysis. (Backward data-flow analysis is similar.) In an ordinary serial CFG, forward data-flow analysis evaluates $\text{IN}(b)$ as the union of $\text{OUT}(a)$ for each predecessor block a of b :

$$\text{IN}(b) = \bigcup_{(a,b) \in E} \text{OUT}(a) .$$

To handle Tapir CFG’s, data-flow analyses must be adapted specifically to handle `reattach` edges. Because Tapir’s asymmetric representation propagates virtual registers and

memory state differently across a reattach edge, the modifications to LLVM data analyses consider registers and memory separately.

For variables stored in shared memory, the standard data-flow equations remain unchanged. Thus, LLVM need not be modified to handle them for Tapir.

For register variables, however, LLVM's data-flow analyses must be modified to exclude the values in registers from an immediate predecessor a of a basic block b if the edge $(a, b) \in E$ is a reattach edge. Denote the set of reattach edges in E by E_R . For a Tapir CFG, forward data-flow analyses define $\text{IN}(b)$ for register variables as

$$\text{IN}(b) = \bigcup_{(a,b) \in E - E_R} \text{OUT}(a),$$

that is, they ignore predecessors across a reattach edge. With this change, Tapir/LLVM correctly propagates register variables through the CFG, never allowing register values in a basic block to use register values set in a logically parallel detached sub-CFG.

Chapter 4

Optimization passes

Tapir enables LLVM’s existing optimization passes [35] to work across parallel control flow. It also enables new optimization passes that specifically target Tapir’s fork-join parallel constructs. This chapter discusses four representative optimizations. Common-subexpression elimination [41, Sec. 12.2] illustrates an optimization pass that “just works” with the additional Tapir instructions. Loop-invariant code motion [41, Sec. 13.2], and tail-recursion elimination [41, Sec. 15.1] were the only two out of LLVM’s roughly 80 optimization passes that required any modification to work effectively on parallel code. Parallel-loop scheduling serves as an example of a new optimization pass.

4.1 Common-subexpression elimination

The common-subexpression elimination (CSE) optimization identifies redundant calculations and transforms the code so that they are only computed once. For example, the expression $(low+high)/2$ in Figure 4-1a is computed in both line 39 and line 40. Tapir/LLVM performs CSE on this code, producing code equivalent to that in Figure 4-1b. Existing mainstream compilers that support fork-join parallelism do not eliminate this common subexpression, however, and they compute $(low+high)/2$ twice. Tapir/LLVM can perform CSE across either a continue edge, as in the example, or a detach edge. Like the vast majority of optimization passes in Tapir/LLVM, CSE “just works” on Tapir code without any modifications to LLVM’s CSE pass.

a

```

36 void search(int low, int high) {
37     if (low == high) search_base(low);
38     else {
39         cilk_spawn search(low, (low+high)/2);
40         search((low+high)/2 + 1, high);
41         cilk_sync;
42     } }

```

b

```

43 void search(int low, int high) {
44     if (low == high) search_base(low);
45     else {
46         int mid = (low+high)/2;
47         cilk_spawn search(low, mid);
48         search(mid + 1, high);
49         cilk_sync;
50     } }

```

Figure 4-1: Example of common-subexpression elimination on a Cilk program. **a** The function `search`, which uses parallel divide-and-conquer to apply the function `search_base` to every integer in the closed interval `[low, high]`. **b** An optimized version of `search`, where the common subexpression `(low+high)/2` in lines 39 and 40 of the original version is computed only once and stored in the variable `mid` in line 46 of the optimized version.

4.2 Loop-invariant code motion

The loop-invariant code motion (LICM) optimization [41, Sec. 13.2] aims to move computations out of loop bodies if they compute the same value on every iteration of the loop. LICM is responsible, for example, for moving the call to `norm` in the parallel loop in Figure 1-1a outside of the loop, as described in Chapter 1. By adapting LICM to handle parallel loops, Tapir/LLVM reduces the asymptotic serial running time of this parallel loop from $\Theta(n^2)$ to $\Theta(n)$.

Tapir/LLVM requires a minor change to LLVM’s LICM pass to handle parallel loops. Consider the CFG illustrated in Figure 2-1, which models the parallel loops in Figure 1-1. For the serial elision of the loop, which would have a similar graph structure except with the continue edge missing, LLVM attempts to find candidate computations to move outside the loop by looking for instructions in the basic blocks of the loop body that dominate the exit block of the loop, such as the block `inc` in Figure 2-1. (The block labeled `exit` is the exit of the function, not the loop exit.) For a parallel loop, however, this analysis fails to identify any code to move due to the existence of the continue edge. As Figure 2-1 shows, with the continue edge, blocks in the loop body can never dominate the exit block `inc` as they could for the serial elision.

Tapir/LLVM modifies LLVM’s LICM pass to handle a parallel loop by analyzing the serial elision of the loop, which essentially means ignoring continue edges. For simple

a

```
51 void pqsort(int* start, int* end) {
52     if (begin == end) return;
53     int* mid = partition(start, end);
54     swap(end, mid);
55     cilk_spawn pqsort(begin, mid);
56     pqsort(mid+1, end);
57     cilk_sync;
58     return;
59 }
```

c

```
60 void pqsort(int* start, int* end) {
61     pqsort_start:
62     if (begin == end) {
63         cilk_sync;
64         return;
65     }
66     int* mid = partition(start, end);
67     swap(end, mid);
68     cilk_spawn pqsort(begin, mid);
69     start = mid+1;
70     goto pqsort_start;
71 }
```

b

```
72 void pqsort(int* start, int* end) {
73     if (begin == end) return;
74     int* mid = partition(start, end);
75     swap(end, mid);
76     cilk_spawn pqsort(begin, mid);
77
78     start = mid+1;
79     // Begin inlined code
80     if (begin == end) goto join;
81     mid = partition(start, end);
82     swap(end, mid);
83     cilk_spawn pqsort(begin, mid);
84     pqsort(mid+1, end);
85     cilk_sync;
86     // End inlined code
87
88     join:
89     cilk_sync;
90     return;
91 }
```

Figure 4-2: Example of tail-recursion elimination on a parallel quicksort program. **a** The Cilk function `pqsort` sorts an array of integers in the range specified by the `start` and `end` pointers. **b** A version of `pqsort` where the recursive tail call on line 56 has been replaced by one round of inlining. **c** A version of `pqsort` where tail-recursion elimination has removed the recursive tail call on line 56.

parallel loop structures with a single continue edge, such as that shown in Figure 2-1, this modification is implemented by finding blocks in the loop body that dominate the predecessors of the loop exit. The modification required changing only 25 lines of LLVM's LICM pass.

4.3 Tail-recursion elimination

Tail-recursion elimination (TRE) [41, Sec. 15.1] aims to replace a recursive call at the end of a function with a branch to the start of the function. By eliminating these recursive tail calls, TRE can avoid function-call overheads and reduce the stack space they consume. This optimization can especially benefit fork-join parallel programs, as many parallel runtime systems impose additional setup and cleanup overhead on a spawned function.

LLVM’s existing TRE pass can perform the TRE optimization on Tapir programs with just a minor modification. Specifically, the modified TRE pass ignores `sync` instructions after the tail-recursive call. Further, if TRE is applied and ignores a `sync` instruction, it must then insert a `sync` instruction before any remaining returns. This modification to LLVM’s TRE pass required changing only 68 lines.

To see why these `sync` instructions can be safely ignored, consider Figure 4-2, which illustrates how Tapir/LLVM’s TRE pass operates on the `pqsort` function, a parallel version of Hoare’s quicksort algorithm [14]. The original tail-recursive code is shown in Figure 4-2a. Figure 4-2b illustrates the result of simply inlining the tail-recursive call. For the inlined code, all `return` statements are replaced with branches to the `join` label. Because there is a `cilk_sync` at the start of `join`, the `cilk_sync` on line 85 can be eliminated. call an arbitrary number of times, TRE can safely ignore a `cilk_sync` instruction after the final tail-recursive call, assuming that it inserts a `cilk_sync` instruction before all remaining returns.

4.4 Parallel-loop scheduling and lowering

As discussed above and in Chapter 2, Tapir effectively represents a parallel loop as a serial loop over a body that is spawned every iteration. Depending on the number of iterations of the loop and the amount of work inside each loop, however, statically scheduling loop iterations in this way may be inefficient. For a parallel loop with a large number of iterations, for instance, it is faster to schedule the iterations in a recursive divide-and-conquer fashion, which produces more parallelism (see [39, Sec. 8.3]). For parallel loops with few iterations, however, the additional function calls required to perform the parallel divide-and-conquer can make the loop run slower than simply spawning off the iterations.

Tapir/LLVM implements a parallel optimization pass that schedules the iterations of a parallel loop using recursive divide-and-conquer, but only if that loop contains sufficiently many iterations. This pass is implemented as part of Tapir/LLVM’s 3800-line lowering pass, which translates `detach`, `reattach`, and `sync` instructions into appropriate Cilk Plus runtime calls [15]. In particular, Tapir/LLVM uses the Cilk Plus runtime calls for `cilk_for`

loops [15, Sec 10.7] to schedule parallel loops. Although we could have separated parallel-loop scheduling from lowering, we chose to combine these two passes so that we could perform fair comparisons between Tapir/LLVM and compilers that lower parallel constructs in their front end. We plan to separate the parallel-loop-scheduling and lowering passes in a future version of Tapir/LLVM.

4.5 Other optimization passes

Tapir/LLVM implements two minor parallel optimization passes: unnecessary-synchronization elimination and puny-task elimination. *Unnecessary-synchronization elimination* identifies and eliminates `sync` instructions that could not possibly sync a detached sub-CFG. *Puny-task elimination* serializes detached sub-CFG's that perform little or no work. If the runtime overhead of creating a parallel task outweighs the work in the task, the task might as well be run serially. Both of these optimization passes were implemented in 52 lines of code by augmenting LLVM's SimplifyCFG pass.

Chapter 5

Auxiliary software

This chapter describes auxiliary software that the Tapir team developed to exercise and test Tapir/LLVM. Although our research focuses on the middle end of the compiler, we implemented a front end for Cilk Plus. In addition, we developed compiler instrumentation that allows the compiler to interface to a race detector to verify the correctness of the Tapir/LLVM implementation.

To create the front end, the Tapir team created a modification of the Clang front end called PClang, which translates Cilk Plus codes to Tapir. We also created a version of Clang that can handle some OpenMP codes. PClang handles most of the fork-join control constructs specified by the Cilk Plus programming model, and specifically, enough to run all the benchmarks described in Chapter 6.

We augmented Tapir/LLVM in two ways to test the correctness of the implementation. First, we modified LLVM's internal verification pass to check that Tapir's invariants are also maintained. Second, we added an instrumentation pass to Tapir/LLVM to allow parallel executables to be tested for determinacy races using a provably good determinacy race detector. This race detector, based on the SP-bags algorithm [11], is guaranteed to find a determinacy race if and only if one exists in the program execution. The verification pass and race detector helped us locate and fix bugs in Tapir/LLVM, both within our code and within the underlying LLVM codebase. Tapir/LLVM now passes all tests in LLVM's regression test suites and correctly compiles our own suite of parallel test programs.

The instrumentation pass has proved useful for supporting other dynamic-analysis tools

based on Tapir/LLVM. Genghis Chau of MIT adapted the Cilkprof scalability profiler [54] to use Tapir/LLVM and this instrumentation in order to build an integrated development environment with always-on race detection and scalability profiling facilities.

Chapter 6

Evaluation

To evaluate the effectiveness of the approach, the Tapir team evaluated Tapir/LLVM on 20 benchmarks. The experiments support the contention that Tapir’s approach of embedding parallelism in the IR is superior to lowering parallelism in the compiler front end. We could not simply run Tapir/LLVM against another compiler, such as Cilk Plus/LLVM [17], which lowers parallelism in the front end, because Cilk Plus/LLVM and Tapir/LLVM differ in more ways than just where they lower parallel constructs. Consequently, to perform an apples-to-apples comparison of these two approaches, we implemented a compiler called “Reference,” which is as close to identical to Tapir/LLVM as we could muster, except for where lowering occurs. Figure 6-1 illustrates the compilation pipelines for Clang/LLVM, Tapir/LLVM, and Reference.

The first pipeline, Clang/LLVM, has the traditional three-phase structure. The Clang front-end takes serial C/C++ code and emits LLVM IR. The `-O3` middle-end optimizes the IR, and the CodeGen back-end lowers LLVM IR to machine code for a particular hardware platform.

The second pipeline shows how Tapir/LLVM is organized. The PClang front end takes parallel Cilk Plus code as input and emits Tapir. The middle-end now consists of three steps: `-O3` optimization, a Lower pass to lower Tapir to LLVM IR, and another pass at `-O3` optimization. The first `-O3` pass performs optimizations on the Tapir representation, the lowering pass translates all the Tapir-specific constructs to LLVM IR, and the second `-O3` pass performs optimizations on the LLVM IR. Finally, the CodeGen back end lowers

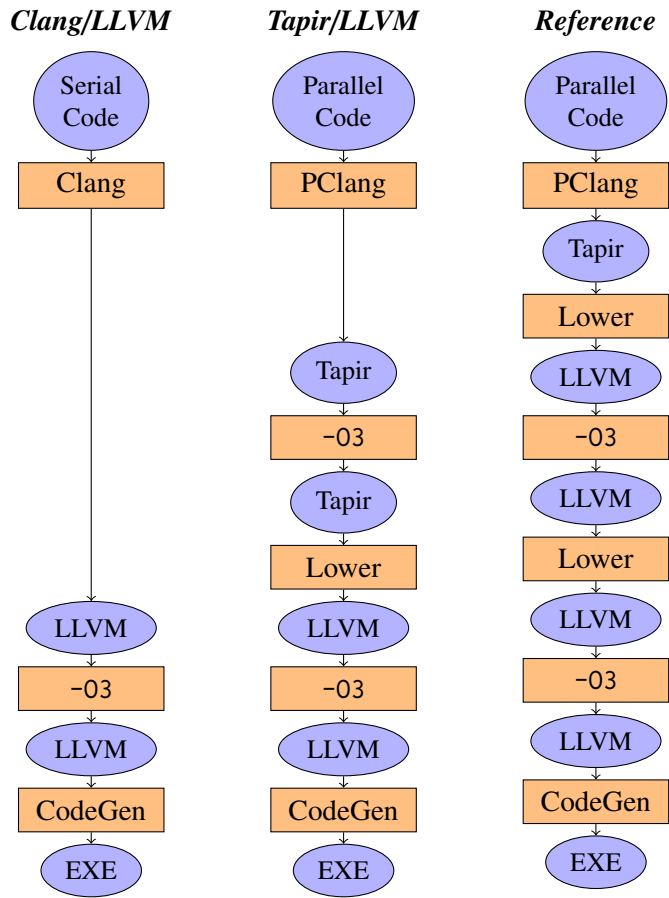


Figure 6-1: The compilation pipelines for Clang/LLVM, Tapir/LLVM, and Reference. Each block represents a compiler transformation, and each oval designates the format of the code at that point in the pipeline.

LLVM IR to machine code.

The third pipeline, called Reference, models how mainstream compilers work today, where parallel constructs are transformed into runtime calls before any optimization can take place. The only difference between Reference and Tapir/LLVM is that the Tapir code emitted by the PClang front end is immediately lowered to LLVM IR before the rest of the Tapir pipeline is invoked. (The second Lower pass in the Reference pipeline therefore has no effect.) Although Reference lowers the parallel constructs early, two iterations of `-O3` are included to ensure that the Tapir/LLVM gains no advantage from optimizing twice. Although one might think that a second pass of `-O3` would be redundant, it is not. For example, a simple matrix-multiplication code runs 13% faster after two rounds of opti-

<i>Suite</i>	<i>Benchmark</i>	<i>Description</i>
<i>Cilk</i>	Cholesky	Cholesky decomposition
	FFT	Fast Fourier transform
	NQueens	n -Queens solver
	QSort	Hoare quicksort
	RectMul	Rectangular matrix multiplication
	Strassen	Strassen matrix multiplication
<i>Intel</i>	AvgFilter	Averaging filter on an image
	Mandel	Mandelbrot set computation
<i>PBBS</i>	CHull	Convex hull
	detBFS	BFS, deterministic algorithm
	incMIS	MIS, incremental algorithm
	incST	Spanning tree, incremental algorithm
	kdTree	Performance test of a parallel k -d tree
	ndBFS	BFS, nondeterministic algorithm
	ndMIS	MIS, nondeterministic algorithm
	ndST	Spanning tree, nondeterministic algorithm
	parallelSF	Spanning-forest computation
	pRange	Compute ranges on a parallel suffix array
	radixSort	Radix sort
	SpMV	Sparse matrix-vector multiplication

Figure 6-2: Descriptions of the 20 benchmarks used to evaluate Tapir/LLVM. These benchmarks were taken from the MIT Cilk benchmark suite [12], Intel Cilk Plus example programs [19], and the CMU Problem-Based Benchmark Suite [57]. “MIS” denotes the computation of a maximal independent set of a graph. “BFS” denotes the breadth-first search of a graph.

mization compared to just one. And although most benchmarks run faster after two `-O3` passes, some actually run slower. Thus, we implemented Reference with the same passes as Tapir/LLVM, except for the initial Lower pass in Reference. This difference only affects parallel code. Serial code passes through both pipelines identically.

6.1 Benchmarking

To benchmark the compiler pipelines, we assembled a collection of benchmark programs taken from the MIT Cilk benchmark suite [12], Intel Cilk code samples [19], and the CMU Problem-Based Benchmark Suite [57]. From these collections, we selected stable programs that tend to exhibit little performance difference when the number or order of optimization

passes is changed. Figure 6-2 describes the suite of benchmarks tested.

We compiled each program in our benchmark suite with both Tapir/LLVM and Reference, and we ran them on both 1 and 18 cores of our test machine. Additionally, we compiled the serial elision of each benchmark with each compiler. Each running time is the minimum of 10 runs on an Amazon AWS c4.8xlarge spot instance, which is a dual-socket Intel Xeon E5-2666 v3 system with a total of 60 GiB of memory. Each Xeon is a 2.9 GHz 18-core CPU with a shared 25 MiB L3-cache. Each core has a 32 KiB private L1-data-cache and a 256 KiB private L2-cache. The system was “quiesced” to permit careful measurements by turning off Turbo Boost, dvfs, hyperthreading, extraneous interrupts, etc.

6.2 Overall performance

The results of our tests are given in Figure 6-3. For the first pair of rows, Reference and Tapir/LLVM produce essentially identical executables when compiling the serial elision of a benchmark. Differences in running times in these rows are due to system noise. The second pair of rows shows that Tapir/LLVM produces executables with better work than Reference on 15 of the benchmarks. Of the remaining 5 benchmarks, 4 demonstrate less than a 1% difference between their work relative to Tapir/LLVM or Reference. The fourth pair of rows elaborates on the results in the second pair to show that Tapir/LLVM produces executables with nearly optimal work efficiency (within 1%) on 12 of the benchmarks, whereas Reference does so on only 2. The third and fifth pairs of row show that Tapir/LLVM generally produces executables with similar or better parallel speedups than those produced by Reference.

The biggest slowdown created from Tapir/LLVM’s compilation occurs on Cholesky, for which the executable produced by Tapir/LLVM has 4% more work than that produced by Reference. In investigating this benchmark, we found that LLVM runs a handful of optimizations on each function before the middle-end optimization and lowering passes in either Tapir/LLVM or Reference. Although these early optimizations have little effect on most programs, they reduce the work of the Reference-compiled Cholesky executable by approximately 20%. Although we experimented with several ways to implement lowering

in Reference before these early optimizations, the resulting compilers consistently exhibited bugs on other benchmarks in the suite. In our final design for Reference, we placed the initial lowering pass as early as we could muster while still ensuring that Reference could compile all benchmarks correctly.

		Cholesky	FFT	NQueens	QSort	RectMul	Strassen	AvgFilter
T_S	Ref.	2.935	10.304	3.084	4.983	10.207	10.105	1.751
	Tapir	2.933	10.271	3.083	4.984	10.207	10.119	1.750
T_1	Ref.	6.581	10.413	10.196	2.355	30.520	1.316	6.596
	Tapir	6.461	10.415	10.196	1.730	25.774	1.187	5.673
T_{18}	Ref.	0.648	0.609	1.106	0.708	1.847	0.124	0.517
	Tapir	0.709	0.611	1.124	0.615	1.559	0.120	0.467
$\frac{T_S}{T_1}$	Ref.	0.757	0.980	0.991	0.743	0.845	0.710	0.801
	Tapir	0.771	0.980	0.991	1.012	1.000	0.788	0.992
$\frac{T_S}{T_{18}}$	Ref.	7.690	16.760	9.137	2.472	13.957	7.540	9.518
	Tapir	7.028	16.705	8.990	2.846	16.536	7.792	10.942
		Mandel	CHull	detBFS	incMIS	incST	kdTree	ndBFS
T_S	Ref.	25.779	0.938	5.670	4.993	4.190	5.473	3.950
	Tapir	25.780	0.935	5.666	5.006	4.173	5.466	3.956
T_1	Ref.	4.572	11.919	3.409	6.030	4.733	5.640	4.930
	Tapir	4.739	11.733	3.419	5.043	4.203	5.546	3.980
T_{18}	Ref.	0.387	0.788	0.196	0.559	0.352	0.342	0.415
	Tapir	0.396	0.774	0.197	0.527	0.329	0.339	0.361
$\frac{T_S}{T_1}$	Ref.	0.642	0.862	0.904	0.828	0.882	0.969	0.801
	Tapir	0.619	0.875	0.902	0.990	0.993	0.986	0.992
$\frac{T_S}{T_{18}}$	Ref.	7.579	13.034	15.730	8.932	11.855	15.982	9.518
	Tapir	7.407	13.270	15.650	9.474	12.684	16.124	10.942
		ndBFS	ndMIS	ndST	parallelSF	pRange	radixSort	SpMV
T_S	Ref.	3.950	9.210	4.069	5.136	2.564	3.775	1.780
	Tapir	3.956	9.253	4.053	5.136	2.559	3.775	1.783
T_1	Ref.	4.930	10.760	4.286	5.646	3.438	3.795	1.836
	Tapir	3.980	9.246	4.063	5.183	3.083	3.800	1.786
T_{18}	Ref.	0.415	0.774	1.925	0.414	0.348	0.284	0.118
	Tapir	0.361	0.701	1.692	0.392	0.330	0.285	0.112
$\frac{T_S}{T_1}$	Ref.	0.801	0.856	0.946	0.910	0.744	0.995	0.969
	Tapir	0.992	0.996	0.998	0.991	0.830	0.993	0.997
$\frac{T_S}{T_{18}}$	Ref.	9.518	11.899	2.105	12.406	7.353	13.292	15.085
	Tapir	10.942	13.138	2.395	13.102	7.755	13.246	15.893

Figure 6-3: Comparison between executables compiled using Reference and using Tapir/LLVM. Each column refers to a different parallel benchmark described in Figure 6-2. Rows labeled “Ref.” describe executables compiled using Reference, and rows labeled “Tapir” describe executables compiled using Tapir/LLVM. Each measured running time is the minimum over 10 executions, measured in seconds. The pair of rows labeled T_S gives the running time of the executable compiled from the serial elision of each benchmark. The pair of rows labeled T_1 gives the work of each benchmark. The pair of rows labeled T_{18} gives the 18-core running time of each benchmark. The pair of rows labeled T_S/T_1 gives the work efficiency of each compiled benchmark, derived from the first and second pairs of rows. The pair of rows labeled T_S/T_{18} gives the parallel speedup of each compiled executable on 18 cores, derived from the first and third pairs of rows.

Chapter 7

Related work

This chapter describes related work in representing parallelism in a compiler IR and in analyzing and optimizing parallel programs.

Various prior research explores compiler optimizations on unstructured parallel threads. For example, some researchers have explored how to find and remove unnecessary synchronization in Java programs [3, 50]. Joisha *et al.* [20] present a technique to detect instructions that are unaffected by parallel threads and can be safely optimized across unstructured parallel control flow. In contrast, our work on Tapir focuses on compiler optimizations for structured parallelism, namely fork-join parallel programs with serial semantics. Although fork-join parallelism may be more restricted than unstructured parallel threads, Tapir demonstrates that many of the optimizations for serial code easily extend to fork-join parallelism. Enabling similar optimizations for unstructured parallel threads appears to be a much harder problem.

Some previous work on compiler optimizations for fork-join parallel programs evaluate which instructions can safely execute in parallel [1] based on concurrency mechanisms supported by a particular memory model. For example, Barik *et al.* [5, 6] use interprocedural analysis to perform various optimizations affecting critical sections of X10 and Habanero-Java programs. Rather than dealing with the complexities of general concurrency mechanisms, Tapir enables compiler optimizations for an easy-to-understand situation: when the optimization respects the serial semantics of the program and does not introduce determinacy races. Compared with general concurrency mechanisms, well-structured parallelism

seems to offer a less onerous path to performance.

Khaldi *et al.* [24] modify LLVM IR to support OpenSHMEM parallel programs with the aim of achieving performance in modern network interconnects that support efficient data transfers for partitioned global address spaces (PGAS). Based on the SPIRE methodology [23] for representing parallel code, they augment functions, basic blocks, instructions, identifiers, and types in LLVM IR with execution, synchronization, scheduling, and memory-layout information. In contrast, Tapir models fork-join parallelism for shared-memory multicores, a conceptually simpler context than PGAS systems, and extends LLVM IR minimally using only three instructions. Once again, the Tapir’s strong assumption of a fork-join programming model with serial semantics that compiles to a flexible multicore architecture seems to provide both performance and simplicity, albeit at the cost of scalability to huge cluster-based supercomputers that lack strong memory-consistency guarantees.

In contrast with much of the work referenced above, Chatarasi *et al.* [9] focus, as Tapir does, on fork-join programs with serial semantics. Specifically, they examine polyhedral optimizations on OpenMP programs with serial semantics. By combining dependency and happens-before analyses, they manage to enable traditional polyhedral optimizers to work on parallel loops, much as Tapir enables common middle-end compiler optimizations to work on parallel code.

Chapter 8

Conclusion

To conclude, I would like to leave the reader with three interesting considerations regarding the nature of asymmetry in parallelism, the future of parallel optimizations, and extensions of Tapir-like systems to other models of parallel programming.

Reasoning about logically parallel tasks asymmetrically based on serial semantics can sometimes simplify the understanding of a parallel program's behavior. When a task is spawned to execute in parallel with another, it is natural to reason about the logically parallel tasks as symmetric, because their instructions can execute in any relative order. For parallel programs with serial semantics, however, it is always valid to execute the program on a single processor, which asymmetrically executes one parallel task to completion before starting the other. Serial semantics encourage an asymmetric representation of parallel control flow that is similar enough to its serial elision that most common analyses and transformations for serial programs work on parallel constructs with little or no modification. In particular, serial semantics enables common optimizations on parallel code that can be invalid under other models of parallelism [64].

One of the great benefits of Tapir is that its strategy for representing parallelism makes it easy to write optimization passes specifically for parallel code. Chapter 4 briefly mentioned some parallel optimization passes we implemented, including parallel-loop scheduling and unnecessary-sync elimination. In addition to helping close the performance gap between serial and parallel versions of code, we hope that the introduction of Tapir will encourage the development and implementation of many more parallel-optimization passes.

Finally, Tapir allows fork-join parallel programs to benefit from both serial and parallel optimizations. Moving forwards, it is natural to wonder whether other models of parallelism, such as pipeline parallelism [27, 42, 10] or data-graph computations [37, 36, 38, 44, 45, 56, 58], can take advantage of the Tapir approach.

Appendix A

Artifact description

This guide describes how to set up Tapir/LLVM and how to download and run our suite of application benchmarks. In particular, this guide focuses on setting up and running three software components:

- the Tapir/LLVM compiler,
- the PClang front end to Tapir/LLVM, and
- the suite of 20 Cilk application benchmarks described in Figure 6-2.

I provide instructions to download and build Tapir/LLVM and PClang. I also provide instructions to download the application benchmark suite and run the Tapir/LLVM compiler on that suite.

We have built and tested Tapir/LLVM, PClang, and the test suite on an x86_64 shared-memory multicore machine running Linux. We provide instructions for obtaining Tapir/LLVM and PClang from my GitHub repositories and setting up the compiler on such a machine. Due to the complexity of the LLVM compiler on which Tapir/LLVM is based, building Tapir/LLVM requires significant computational resources: approximately 50 GiB of disk, 12 GiB of RAM, and anywhere from a few minutes to a couple of hours, depending on the machine. We also provide instructions for obtaining a copy of our test suite from a tarball.

A.1 Building Tapir/LLVM from source

This section describes how to download the source code for Tapir/LLVM and PClang from GitHub and build them. These instructions assume you are building Tapir/LLVM on an x86_64 system running Linux.

System requirements. Building Tapir/LLVM and PClang involves building the LLVM and Clang systems that they extend. Because of the size of the underlying LLVM and Clang codebases, you need a relatively powerful machine in order to build the compiler in a timely fashion. Approximately 50 GiB of disk space and 12 GiB of memory are needed to compile LLVM and Clang. A fresh build of LLVM and Clang can take substantial time to complete, e.g., approximately an hour on one processor of an AWS c4.8xlarge instance. The build script will attempt to use parallel processors to speed up compilation. See <http://llvm.org/docs/CMake.html> for more information on building LLVM and Clang.

1. Install the requisite software to build Tapir/LLVM and PClang, namely, `cmake`, `gcc`, and `git`.
2. Download the sources of Tapir/LLVM and PClang from GitHub:

```
$ git clone --recursive https://github.com/wsmoses/Tapir-Meta.git
```

The source is approximately 800 MiB in size.

3. Compile Tapir/LLVM and PClang:

```
$ cd Tapir-Meta/  
$ bash ./build.sh
```

This script will build Tapir/LLVM and PClang and store the compiled binaries in `Tapir-Meta/tapir/build`. If the build succeeds, the final line of output will be `Installation successful`.

4. Set up your environment variables to use Tapir/LLVM and PClang:

```
$ source ./setup-env.sh
```

This script will add the `Tapir-Meta/tapir/build/bin/` subdirectory to your path, so that the `clang` command will refer to Tapir/LLVM and PClang.

A.2 Running the benchmark suite

This section describes how you can download the application benchmark suite described in Figure 6-2 and test Tapir/LLVM on these benchmarks.

1. Install the requisite software to download and run the tests, namely, `bc`, `libcilkrts`, `numactl`, `python`, `taskset`, and `wget`.
2. Download the tarball containing the application benchmark suite and unpack it:

```
$ wget http://tinyurl.com/TapirLLVMTesting -O testing.tar  
$ tar -xvf testing.tar
```

This tarball is approximately 12 GiB in size. Unpacking the tarball creates the `testing/` subdirectory of the current working directory that contains the application benchmark suite.

3. Run the test script:

```
$ cd testing  
$ ./test.sh
```

The test script takes approximately 7 hours to run. The script compiles each benchmark in the test suite twice using Tapir/LLVM: once as a parallel program, and once as the program's serial elision. All compilations use optimization level `-O3`. The test script runs each compiled executable 10 times using 1 worker thread and 10 times using 18 worker threads.

A.3 Evaluation and expected result

Once the test script finishes running, the results can be summarized into a table similar to Figure 6-3 as follows:

```
$ ./results.sh > results.csv
```

This command will produce `results.csv`, a table of tab-separated values that contains the minimum running time from each set of 10 runs of a particular executable on a particular worker count. The table also contains derived work-efficiency and parallel speedup values for each benchmark program.

Because these results are performance measurements, they are likely to vary from run to run and from system to system. Moreover, the Tapir team is continuing to develop the Tapir/LLVM compiler and PClang, meaning that your results will not precisely match those in Figure 6-3.

You can write your own programs and compile them using PClang and Tapir/LLVM. The PClang front end is *not* a fully featured Cilk front end, however. For more information on the source language parsed by PClang, please see `testing/PClang-README.txt`.

Bibliography

- [1] Shivali Agarwal, Rajkishore Barik, Vivek Sarkar, and Rudrapatna K. Shyamasundar. May-happen-in-parallel analysis of X10 programs. In *PPoPP*, pages 183–193, 2007.
- [2] Alfred V. Aho, Monica S. Lam, Ravi Sethi, and Jeffrey D. Ullman. *Compilers: Principles, Techniques, and Tools*. Addison-Wesley, second edition, 2006.
- [3] Jonathan Aldrich, Craig Chambers, EminGun Sirer, and Susan Eggers. Static analyses for eliminating unnecessary synchronization from Java programs. In Agostino Cortesi and Gilberto Filé, editors, *Static Analysis*, volume 1694 of *Lecture Notes in Computer Science*, pages 19–38. 1999.
- [4] E. Ayguade, N. Coptý, A. Duran, J. Hoefflinger, Yuan Lin, F. Massaioli, X. Teruel, P. Unnikrishnan, and Guansong Zhang. The design of OpenMP tasks. *IEEE Transactions on Parallel and Distributed Systems*, 20(3):404–418, 2009.
- [5] Rajkishore Barik and Vivek Sarkar. Interprocedural load elimination for dynamic optimization of parallel programs. In *PACT*, pages 41–52, 2009.
- [6] Rajkishore Barik, Jisheng Zhao, and Vivek Sarkar. Interprocedural strength reduction of critical sections in explicitly-parallel programs. In *PACT*, pages 29–40, 2013.
- [7] Hans-J. Boehm and Sarita V. Adve. Foundations of the C++ concurrency memory model. In *PLDI*, pages 68–78, 2008.
- [8] Vincent Cavé, Jisheng Zhao, Jun Shirako, and Vivek Sarkar. Habanero-Java: the new adventures of old X10. In *PPPJ*, pages 51–61, 2011.
- [9] Prasanth Chatarasi, Jun Shirako, and Vivek Sarkar. Polyhedral optimizations of explicitly parallel programs. In *PACT*, pages 213–226, 2015.
- [10] Wei Du, Renato Ferreira, and Gagan Agrawal. Compiler support for exploiting coarse-grained pipelined parallelism. In *SC*, pages 8–21, 2003.
- [11] Mingdong Feng and Charles E. Leiserson. Efficient detection of determinacy races in Cilk programs. *Theory of Computing Systems*, 32(3):301–326, 1999.
- [12] Matteo Frigo, Charles E. Leiserson, and Keith H. Randall. The implementation of the Cilk-5 multithreaded language. In *PLDI*, pages 212–223, 1998.

- [13] Dirk Grunwald and Harini Srinivasan. Data flow equations for explicitly parallel programs. In *PPoPP*, pages 159–168, 1993.
- [14] C. A. R. Hoare. Algorithm 64: Quicksort. *CACM*, 4(7):321, 1961.
- [15] Intel Corporation. *Intel Cilk Plus Application Binary Interface Specification*, 2010. Document Number: 324512-001US. Available from https://software.intel.com/sites/products/cilk-plus/cilk_plus_abi.pdf.
- [16] Intel Corporation. *Intel Cilk Plus Language Specification*, 2010. Document Number: 324396-001US. Available from http://software.intel.com/sites/products/cilk-plus/cilk_plus_language_specification.pdf.
- [17] Intel Corporation. Cilk Plus/LLVM. Available from <http://cilkplus.github.io/>, 2013.
- [18] Intel Corporation. *Intel C++ Compiler 16.0 User and Reference Guide*, 2015.
- [19] Intel Corporation. Intel Cilk Plus samples. Available from <https://software.intel.com/en-us/code-samples/intel-compiler/intel-compiler-features/intelcilkplus>, 2016.
- [20] Pramod G. Joisha, Robert S. Schreiber, Prithviraj Banerjee, Hans J. Boehm, and Dhruva R. Chakrabarti. A technique for the effective and automatic reuse of classical compiler optimizations on multithreaded code. In *POPL*, pages 623–636, 2011.
- [21] Herbert Jordan, Simone Pellegrini, Peter Thoman, Klaus Kofler, and Thomas Fahringer. INSPIRE: The Insieme parallel intermediate representation. In *PACT*, pages 7–18, 2013.
- [22] Brian W. Kernighan and Dennis M. Ritchie. *The C Programming Language*. Prentice Hall, Inc., second edition, 1988.
- [23] D. Khaldi, P. Jouvelot, C. Ancourt, and F. Irigoin. SPIRE, a sequential to parallel intermediate representation extension. Technical report, Technical Report CRI/A-487, MINES ParisTech, 2012.
- [24] Dounia Khaldi, Pierre Jouvelot, François Irigoin, Corinne Ancourt, and Barbara Chapman. LLVM parallel intermediate representation: Design and evaluation using OpenSHMEM communications. In *LLVM*, pages 2:1–2:8, 2015.
- [25] Jens Knoop, Bernhard Steffen, and Jürgen Vollmer. Parallelism for free: Efficient and optimal bitvector analyses for parallel programs. *ACM TOPLAS*, pages 268–299, 1996.
- [26] Chris Lattner and Vikram Adve. LLVM: A Compilation Framework for Lifelong Program Analysis & Transformation. In *CGO*, pages 75–87, 2004.
- [27] I-Ting Angelina Lee, Charles E. Leiserson, Tao B. Schardl, Zhunping Zhang, and Jim Sukha. On-the-fly pipeline parallelism. *ACM TOPC*, 2(3):17:1–17:42, 2015.

- [28] Jaejin Lee, Samuel P. Midkiff, and David A. Padua. Concurrent static single assignment form and constant propagation for explicitly parallel programs. In *LCPC*, pages 114–130, 1997.
- [29] Charles E. Leiserson. The Cilk++ concurrency platform. *Journal of Supercomputing*, 51(3):244–257, 2010.
- [30] LLVM Developer List. LLVMdev discussions on Intel OpenMP proposal. Available from <http://lists.llvm.org/pipermail/llvm-dev/2012-September/053861.html>, September 2012.
- [31] LLVM Developer List. LLVMdev discussions on OpenCL SPIR proposal. Available from <http://lists.llvm.org/pipermail/llvm-dev/2012-September/053293.html>, September 2012.
- [32] LLVM Developer List. LLVMdev Parallelization metadata and intrinsics in LLVM (for OpenMP, etc.). Available from <http://lists.llvm.org/pipermail/llvm-dev/2012-September/053792.html>, September 2012.
- [33] LLVM Developer List. LLVMdev discussions on parallel IR. Available from <http://lists.llvm.org/pipermail/llvm-dev/2015-March/083314.html>, March 2015.
- [34] LLVM Project. *LLVM Language Reference Manual*, 2015. Available from <http://llvm.org/docs/LangRef.html>.
- [35] LLVM Project. *LLVM’s Analysis and Transform Passes*, 2015. Available from <http://llvm.org/docs/Passes.html>.
- [36] Yucheng Low, Danny Bickson, Joseph Gonzalez, Carlos Guestrin, Aapo Kyrola, and Joseph M. Hellerstein. Distributed GraphLab: a framework for machine learning and data mining in the cloud. *Proceedings of the VLDB Endowment*, pages 716–727, 2012.
- [37] Yucheng Low, Joseph Gonzalez, Aapo Kyrola, Danny Bickson, Carlos Guestrin, and Joseph M. Hellerstein. GraphLab: A new framework for parallel machine learning. In *UAI*, pages 340–349, 2010.
- [38] Grzegorz Malewicz, Matthew H. Austern, Aart J.C Bik, James C. Dehnert, Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. Pregel: a system for large-scale graph processing. In *SIGMOD*, pages 135–146, 2010.
- [39] Michael McCool, Arch D. Robison, and James Reinders. *Structured Parallel Programming: Patterns for Efficient Computation*. Elsevier Science, 2012.
- [40] Samuel P. Midkiff and David A. Padua. Issues in the optimization of parallel programs. In *ICPP*, pages 105–113, 1990.
- [41] Steven S. Muchnick. *Advanced Compiler Design and Implementation*. Morgan Kaufmann, 1997.

- [42] Angeles Navarro, Rafael Asenjo, Siham Tabik, and Calin Cascaval. Analytical modeling of pipeline parallelism. In *PACT*, pages 281–290, 2009.
- [43] Robert H. B. Netzer and Barton P. Miller. What are race conditions? *ACM Letters on Programming Languages and Systems*, 1(1):74–88, 1992.
- [44] Donald Nguyen, Andrew Lenharth, and Keshav Pingali. A lightweight infrastructure for graph analytics. In *SOSP*, pages 456–471, 2013.
- [45] Donald Nguyen, Andrew Lenharth, and Keshav Pingali. Deterministic Galois: On-demand, portable and parameterless. In *ASPLOS*, pages 499–512, 2014.
- [46] Diego Novillo, Ron Unrau, and Jonathan Schaeffer. Concurrent SSA form in the presence of mutual exclusion. In *ICPP*, pages 356–364, 1998.
- [47] OpenMP Architecture Review Board. *OpenMP Application Program Interface, Version 4.0*, July 2013. Available from <http://www.openmp.org/mp-documents/OpenMP4.0.0.pdf>.
- [48] Antoniu Pop and Albert Cohen. Preserving high-level semantics of parallel programming annotations through the compilation flow of optimizing compilers. In *CPC*, 2010.
- [49] William Pugh. Fixing the Java memory model. In *JAVA*, pages 89–98, 1999.
- [50] Erik Ruf. Effective synchronization removal for Java. In *PLDI*, pages 208–218, 2000.
- [51] Radu Rugina and Martin C. Rinard. Pointer analysis for structured parallel programs. *TOPLAS*, pages 70–116, January 2003.
- [52] Vivek Sarkar. Analysis and optimization of explicitly parallel programs using the parallel program graph representation. In *LCPC*, pages 94–113, 1998.
- [53] Vivek Sarkar and Barbara Simons. Parallel program graphs and their classification. In *LCPC*, pages 633–655, 1994.
- [54] Tao B. Schardl, Bradley C. Kuszmaul, I-Ting Angelina Lee, William M. Leiserson, and Charles E. Leiserson. The Cilkprof scalability profiler. In *SPAA*, pages 89–100, 2015.
- [55] Tao B. Schardl, William S. Moses, and Charles E. Leiserson. Tapir: Embedding fork-join parallelism into LLVM’s intermediate representation. In *Proc. 22nd Symposium on Principles and Practice of Parallel Programming*, 2017.
- [56] Julian Shun and Guy E. Blelloch. Ligra: A lightweight graph processing framework for shared memory. In *PPoPP*, pages 135–146, 2013.
- [57] Julian Shun, Guy E. Blelloch, Jeremy T. Fineman, Phillip B. Gibbons, Aapo Kyrola, Harsha Vardhan Simhadri, and Kanat Tangwongsan. Brief announcement: the Problem Based Benchmark Suite. In *SPAA*, pages 68–70, 2012.

- [58] Julian Shun, Laxman Dhulipala, and Guy E. Blelloch. Smaller and faster: Parallel processing of compressed graphs with Ligra+. In *DCC*, pages 403–412, 2015.
- [59] Harini Srinivasan and Dirk Grunwald. An efficient construction of parallel static single assignment form for structured parallel programs. Technical report, Technical Report CU-CS-564-91, University of Colorado at Boulder, 1991.
- [60] Harini Srinivasan, James Hook, and Michael Wolfe. Static single assignment for explicitly parallel programs. In *POPL*, pages 260–272, 1993.
- [61] Harini Srinivasan and Michael Wolfe. Analyzing programs with explicit parallelism. In *LCPC*, pages 405–419, 1991.
- [62] Richard M. Stallman and the GCC Developer Community. *Using the GNU Compiler Collection (for GCC version 6.1.0)*. Free Software Foundation, 2016.
- [63] Bjarne Stroustrup. *The C++ Programming Language*. Addison-Wesley, 4th edition, 2013.
- [64] Viktor Vafeiadis, Thibaut Balabonski, Soham Chakraborty, Robin Morisset, and Francesco Zappa Nardelli. Common compiler optimisations are invalid in the C11 memory model and what we can do about it. In *POPL*, pages 209–220, 2015.
- [65] Jisheng Zhao and Vivek Sarkar. Intermediate language extensions for parallelism. In *SPLASH*, pages 329–340, 2011.