

The Essence of Compiling with Traces

Shu-yu Guo Jens Palsberg

UCLA Computer Science Department
University of California, Los Angeles, USA
{shu,palsberg}@cs.ucla.edu

Abstract

The technique of trace-based just-in-time compilation was introduced by Bala et al. and was further developed by Gal et al. It currently enjoys success in Mozilla Firefox’s JavaScript engine. A trace-based JIT compiler leverages run-time profiling to optimize frequently-executed paths while enabling the optimized code to “bail out” to the original code when the path has been invalidated. This optimization strategy differs from those of other JIT compilers and opens the question of *which trace optimizations are sound*. In this paper we present a framework for reasoning about the soundness of trace optimizations, and we show that some traditional optimization techniques are sound when used in a trace compiler while others are unsound. The converse is also true: some trace optimizations are sound when used in a traditional compiler while others are unsound. So, traditional and trace optimizations form incomparable sets. Our setting is an imperative calculus for which tracing is explicitly spelled out in the semantics. We define optimization soundness via a notion of bisimulation, and we show that sound optimizations lead to confluence and determinacy of stores.

Categories and Subject Descriptors D.2.4 [Program Verification]: Correctness proofs, formal methods; D.3.4 [Processors]: Compilers; F.3.2 [Semantics of Programming Languages]: Operational semantics

General Terms Languages, Theory

Keywords just-in-time compilation, compiler correctness, bisimulation

1. Introduction

With the advent of “Web 2.0”, the web browser has become a platform that delivers rich interactive applications. The technology central to this transformation of the web browser is JavaScript. JavaScript’s dynamic nature has since then become a performance bottleneck. The performance of dynamic languages is much worse than statically typed languages, and JavaScript is no exception. Moreover, traditional just-in-time (JIT) compilation techniques designed for static, typed languages are ill-fitted for JavaScript.

The work of Bala et al. [1] was adapted as a novel JIT compilation technique called trace compilation [2–4, 7, 8]. A trace-based JIT compiler uses run-time profiling to approximate the “hot” exe-

cution paths (loops) in the program and compiles only those paths [4, 7, 8]. The rarely executed bits of code are interpreted. The idea is quite intuitive: if there is a repeatedly executed section of the code, that section should be top priority for compiling to native code. For example, a micro-blogging web application might take many rows of data and transform them into a news feed format. This loop would be where the program spends the majority of its time; a trace-enabled JIT detects that this loop is a hot execution path and compiles it to native code.

Tracing JIT compilers are amenable to JavaScript and enjoy the greatest success in Mozilla Firefox’s TraceMonkey JavaScript engine. It is available in versions 3.5 and greater, and Mozilla metrics report that approximately 94 million people in the world are using the tracing JIT [13].¹

Tracing JIT compilers differ greatly in technique from many other JIT techniques. It opens the following question:

Which trace optimizations are sound?

We distill the essence of trace compilation to a simple imperative calculus with an operational semantics. This allows us to formally investigate notions of correctness of trace-based JIT compilers and the properties that trace optimizations must satisfy to be sound. We present a bisimulation-based soundness criterion for trace optimizations, and we prove a determinism theorem: whether one traces or not, the final store will be the same.

Our framework is modular in two ways. First, an optimization designer needs only prove that a given optimization satisfies our correctness criterion; the determinism theorem then follows. Second, the composition of two sound optimizations is itself sound. We leverage the first kind of modularity to easily prove soundness of the folding of free loop variables and dead branch elimination. Proving optimizations unsound is equally simple. We show that dead store elimination is unsound with an easy-to-check counterexample. Readers can easily proceed like we did to prove additional trace optimizations sound.

Our proof of the determinism theorem has the following coarse steps. First we prove that an unoptimized, recorded trace of the loop is “behaviorally correct”, or bisimilar, to the original loop. We then prove that the original program with the new trace stitched in place of the old loop is bisimilar to the original program. This then sets the stage for sound optimizations: sound optimizations are those that do not invalidate this behavioral correctness guarantee had from bisimilarity. Finally, we put the pieces together and prove confluence and determinacy of stores via a diamond lemma and a strip lemma.

Our framework shows, surprisingly, that “traditional” whole-function optimizations and tracing JIT optimizations *do not stand in a subset relation in either direction*. In one direction, it is clear

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

POPL’11, January 26–28, 2011, Austin, Texas, USA.
Copyright © 2011 ACM 978-1-4503-0490-0/11/01...\$10.00

¹ The exact average of daily usage from January 1, 2010 until July 12, 2010 of versions 3.5, 3.6, 3.7, and 4.0 is 93,977,941.

that traditional optimizations are not subsumed by trace optimizations. Informally, trace optimizations are not obliged to be correct for all possible executions and contexts. They are obliged to be correct only for a *particular* execution and a *particular* context. For instance, in the paper we prove folding of variables that are not assigned to be sound for tracing but unsound in general in a traditional setting. In the other direction, trace optimizations are also not subsumed by traditional optimizations. The reason for this is more subtle: the domain of trace optimizations is restricted to *only* the trace. The code surrounding the trace is unavailable to the optimizer. Traditional optimizations, on the other hand, are privy to both the prefix and the suffix of the trace in that their domain *is* the entire procedure. In other words, those optimizations can prove properties on entire procedures while trace optimizations cannot. For one, whole-function optimizations know that their local variables are dead after the function exits. Trace optimizations cannot make the same assumption about their local variables after the trace exits. In the paper we show dead store elimination to be unsound as a trace optimization.

To expand upon the incomparability of the two sets of optimizations, it is illuminating to spell out the differences between our work and recent prominent works in compiler correctness [11, 12]. The program equivalence condition (Lemma 1) found in Lacey et al. [11] basically states classical bisimilarity as the condition under which to judge the correctness of optimizations. An optimized program must be bisimilar to the original and their respective final stores must contain the same values for all variables. In their framework, optimizations are formulated as rewrite rules with side conditions expressed in temporal logic. Despite this difference, we can make the following fruitful comparisons. We write $m \approx n$ for σ to mean that the program m is bisimilar to n when their initial stores are σ . Suppose there is an optimization function O . Their correctness criterion describes traditional optimizations and is extensional: for all p , $p \approx O(p)$ for all stores. Our correctness criterion describes tracing optimizations and is intensional. Suppose the program p decomposes into two components, w (the traced loop) and k (the rest of the program). The criterion is then: for all w, k, σ , $w k \approx O(w, \sigma) k$ for σ . Note how our optimization function takes a state σ in addition to a program to produce an optimized program guaranteed to be correct when the initial store is fixed to be that store and the rest of the program is fixed to be k . This succinctly captures that trace optimizations may not be straightforwardly used in a classical setting.

The correctness criterion in Chambers et al. [12] raises yet more differences between classical optimization soundness and trace optimization soundness. At the heart of their formulation of soundness is contextual equivalence. However, note that we have said that trace optimization correctness is intensionalized to a particular computation suffix. This necessarily precludes trace optimization correctness from being contextual equivalence. Classical optimizations speculate about the behavior of a program to substitute optimized portions for the original portions *before* execution. Trace optimizations, on the other hand, know exactly the behavior for which they need to optimize and substitute optimized portions *during* execution. This departure also highlights that the input to trace optimizations is mercurial: we do not know *a priori* what loops are hot. The input to classical optimization on the other hand is fixed. In Chambers et al. and Lacey et al. [11, 12], this input is the entire program. As mentioned before, the domains of the two kinds of optimizations are simply different.

The remainder of this paper is organized as follows. In section 2 we introduce our language, its operational semantics, and discuss certain properties of compiling with traces. In section 3 we show that the operational semantics of our language is correct up to weak bisimulation and relate the results to confluence. In section 4 we

explore various optimizations, both provably correct and provably incorrect, pluggable into the framework. In section 5 we discuss related work. Section 6 concludes.

2. Compiling with Traces

What is essential to the trace compilation technique? What features must our calculus contain? At its core, it is a method of compiling often-executed loops. We must therefore have loops. More specifically, it is a technique of recording loop bodies at run-time and optimizing them. The second essential part must thus be the ability to record execution. What is optimized, then, is not the text of the loop but a run-time execution path *through* the loop. In other words, we are optimizing some fixed execution. So, when the execution diverges from the recorded path, there must exist a mechanism to return us to the original program. The third and final component is this bail-out mechanism. In the literature of trace compilation, these are called side-exits [8].

We forego modeling the many other features of the technique that exist in implementations. For instance, we shall not model the heuristics in how one actually detects a hot path of execution, we will simply build in the ability to record a path of execution nondeterministically. This nondeterminism will be realized via overlapping reduction rules. Nor shall we model implementation details, such as trace trees and their interactions [7]. We aim to keep the calculus minimal yet high-level, safe—and even desirable—for human consumption.

2.1 The Language and Its Baseline Execution

We first present the syntax and small-step operational semantics for a simple imperative language with two of the three essential ingredients: loops and the ability to “bail out” modeled by continuations.

The syntax of our language, inspired by the calculus presented in Moll [9], is shown in figure 1. Concretely, traces are a subset of normal statements. They are meant to be straightline sections of code with side-exits, so there are only no-ops, assignments, and side-exits.

We use x to range over variables, i to range over non-negative integers, σ, ρ to range over stores, and l, m, n, k, s, p to range over statements throughout the rest of the paper.

The baseline transition rules are shown in figure 2. Assume \oplus is the “real” addition operator on integers. We use a labeled transition system where labels correspond to store updates. We assume the reader is familiar with such systems as they are used in the literature of concurrency [14]. The only observable transitions are store updates, which are labeled by the “store delta”. All other transitions are silent, labeled τ , and are unobservable. The subscript B denotes baseline transition rules. The subscript A denotes a strict subset of the baseline transition rules that will be used in the upcoming proofs. The subscript T denotes tracing transition rules. The baseline rules in figure 2 are common to both.

The baseline rules do the usual things. `BailTrue` is the rule that applies continuations in the **bails**. It says to clobber the current reduct with the packaged continuation s .

2.2 Recording Traces

We extend the baseline execution with the ability to record traces. The set of baseline rules is a proper subset of the tracing rules, i.e. $\rightarrow_{BC} \rightarrow_T$. The abstract syntax is the same between the two languages. The additional transition rules are shown in figure 3.

Starting Traces We start a trace at the beginning of a **while** loop. For technical reasons for the proof of correctness, we record when we have already unfolded at least one iteration of the loop.

Also note that Trace puts the reduction rules in *recording mode*, which is represented syntactically as 4-tuples. The components are,

$e ::= n \mid x + 1$	expressions
$b ::= x = 0 \mid x \neq 0$	boolean expressions
$w ::= \mathbf{while} \ b \ \mathbf{do} \ s$	loops
$s ::= \epsilon \mid c \ s$	statements
$c ::= \mathbf{skip}; \mid x := e; \mid w \mid \mathbf{if} \ b \ \mathbf{then} \ s \mid \mathbf{bail} \ b \ \mathbf{to} \ s$	commands
$t ::= \epsilon \mid c_t \ t$	traces
$c_t ::= \mathbf{skip}; \mid x := e; \mid \mathbf{bail} \ b \ \mathbf{to} \ s$	recorded commands

Figure 1. Syntax of the Simple Imperative Language and Traces

$\hat{\sigma}(e) = \begin{cases} n & \text{if } e = n \\ \sigma(x) \oplus 1 & \text{if } e = x + 1 \end{cases}$	$\hat{\sigma}(b) = \begin{cases} \mathit{true} & \text{if } b \text{ is } x = 0 \wedge \sigma(x) = 0 \\ \mathit{false} & \text{if } b \text{ is } x = 0 \wedge \sigma(x) \neq 0 \\ \mathit{true} & \text{if } b \text{ is } x \neq 0 \wedge \sigma(x) \neq 0 \\ \mathit{false} & \text{if } b \text{ is } x \neq 0 \wedge \sigma(x) = 0 \end{cases}$
$\delta ::= x/i \mid x/\mathit{true} \mid x/\mathit{false}$	store updates
$\alpha ::= \tau \mid \delta$	actions
$\langle \sigma, x := e; k \rangle \xrightarrow{\delta}_{T,B,A} \langle \sigma[x/\hat{\sigma}(e)], k \rangle$	(Assign)
$\langle \sigma, \mathbf{skip}; k \rangle \xrightarrow{\tau}_{T,B,A} \langle \sigma, k \rangle$	(Seq)
$\langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ s) \ k \rangle \xrightarrow{\tau}_{T,B,A} \langle \sigma, k \rangle$	(IfFalse)
$\langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ s) \ k \rangle \xrightarrow{\tau}_{T,B,A} \langle \sigma, s \ k \rangle$	(IfTrue)
$\langle \sigma, (\mathbf{while} \ b \ \mathbf{do} \ s) \ k \rangle \xrightarrow{\tau}_{T,B,A} \langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s \ \mathbf{while} \ b \ \mathbf{do} \ s)) \ k \rangle$	(While)
$\langle \sigma, (\mathbf{bail} \ b \ \mathbf{to} \ s) \ k \rangle \xrightarrow{\tau}_{T,B,A} \langle \sigma, k \rangle$	(BailFalse)
$\langle \sigma, (\mathbf{bail} \ b \ \mathbf{to} \ s) \ k \rangle \xrightarrow{\tau}_{T,B} \langle \sigma, s \rangle$	(BailTrue)

Figure 2. Shared Transition Rules

$\neg b = \begin{cases} x = 0 & \text{if } b \text{ is } x \neq 0 \\ x \neq 0 & \text{if } b \text{ is } x = 0 \end{cases}$	
$\langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s \ (\mathbf{while} \ b \ \mathbf{do} \ s))) \ k \rangle \xrightarrow{\tau}_T \langle \sigma, (\mathbf{while} \ b \ \mathbf{do} \ s) \ k, \epsilon, s \ (\mathbf{while} \ b \ \mathbf{do} \ s) \ k \rangle$	(Trace)
$\langle \sigma, k_w, t, x := e; k \rangle \xrightarrow{\delta}_T \langle \sigma[x/\hat{\sigma}(e)], k_w, t (x := e); k \rangle$	(RecordAssign)
$\langle \sigma, k_w, t, \mathbf{skip}; k \rangle \xrightarrow{\tau}_T \langle \sigma, k_w, t (\mathbf{skip}); k \rangle$	(RecordSeq)
$\langle \sigma, k_w, t, (\mathbf{if} \ b \ \mathbf{then} \ s) \ k \rangle \xrightarrow{\tau}_T \langle \sigma, k_w, t (\mathbf{bail} \ b \ \mathbf{to} \ (s \ k)), k \rangle$	(RecordIfFalse)
$\langle \sigma, k_w, t, (\mathbf{if} \ b \ \mathbf{then} \ s) \ k \rangle \xrightarrow{\tau}_T \langle \sigma, k_w, t (\mathbf{bail} \ \neg b \ \mathbf{to} \ k), s \ k \rangle$	(RecordIfTrue)
$\langle \sigma, k_w, t, (\mathbf{while} \ b \ \mathbf{do} \ s) \ k \rangle \xrightarrow{\tau}_T \langle \sigma, k_w, t (\mathbf{skip}); (\mathbf{if} \ b \ \mathbf{then} \ (s \ \mathbf{while} \ b \ \mathbf{do} \ s)) \ k \rangle$	(RecordWhile)
$\langle \sigma, k_w, t, (\mathbf{while} \ b \ \mathbf{do} \ s) \ k \rangle \xrightarrow{\tau}_T \langle \sigma, O(\mathbf{while} \ b \ \mathbf{do} \ t, \sigma) \ k \rangle$	(Stitch)
$\langle \sigma, k_w, t, k \rangle \xrightarrow{\alpha}_T \langle \sigma', k' \rangle$	(Abort)
	$\text{if } \langle \sigma, k \rangle \xrightarrow{\alpha}_T \langle \sigma', k' \rangle \wedge k_w \neq k$

Figure 3. Tracing Transition Rules

in order, the store, the stopping point of the trace, the trace thus far, and the current program being reduced.

Recording Traces The recording rules record one command at a time and concatenate it to the end of the trace. Concatenation is simple juxtaposition. The trace itself is a straightline section of code, so we install side-exits (pieces of code that jump back to untraced code when the condition we traced no longer holds) when we record conditionals.

To ease the task of proving correctness, RecordWhile appends a **skip** to the trace while unrolling the loop. Its side condition is to ensure that we are recording an *inner* loop inside the current loop we are tracing, and that we have not come full circle and finished tracing. The work for finishing up a trace is done in Stitch, whose side condition is mutually exclusive with that of RecordWhile.

Ending Well-Behaved Traces We end the trace and stitch it back into the program using Stitch when we finish tracing the body of the loop. We know we have finished when we come back to reducing the same loop that started the trace.

We “compile” the loop that was traced into the same language.² The actual optimization is immaterial to the semantics; we assume that there is a sound optimization function $O : (Statement \times Store) \rightarrow Statement$. What soundness entails here will be made precise when we investigate correctness. Informally, soundness means that the output of the O function “does the same thing” as the original code, as far as observable behavior (store updates) goes.

Ending Badly-Behaved Traces We are not guaranteed to finish tracing the body of the loop. That loop might never terminate! Consider the following example; assume s_2 never changes b to 0.

```

1  a := 1;
2  b := 1;
3  while a ≠ 0 do
4      s1
5      while b ≠ 0 do
6          s2

```

If we start tracing the *outer* loop, once we start executing the *inner* loop we will *never finish the outer loop body*, and thus never finish tracing. Implementations of trace compilation, then, must use heuristics to end the trace if it is continuing on for too long.

In our semantics, we model this by introducing another non-deterministic rule that prematurely stops the trace, Abort. This rule shares the same premises with *all* Record_ rules, where “_” is a wildcard. Note that there are no axioms for recording **bails**—what this means is that instead of the semantics getting stuck when trying to trace a trace, we abort the trace (that is, we do not model higher-order tracing). Also note that Abort’s³ side condition is mutually exclusive with Stitch, which is intuitively the “good” situation of a successful trace. In this way the rule models the semantics of bailing out of tracing mode for all “bad” situations.

2.3 Example Trace Recording

To help illustrate the tracing rules and to build some concrete intuition, consider the following contrived example.

Example Input

```

1  x := 0;
2  while x = 0 do
3      y := 0;

```

²Note that this is a simplification in our model. In actual tracing JITs, the compiled code is in machine language.

³The rule is modeled as it is instead of the viable alternative of $\langle \sigma, k_w, t, k \rangle \xrightarrow{\tau} \langle \sigma, k \rangle$ if $k_w \neq k$ for a cleaner proof of correctness.

```

4  while y = 0 do
5      y := 1;
6      z := 1;
7      b := a + 1;

```

There are two loops; the inner loop only iterates once. The variable a is computed at some earlier point in the program. We give a rule-by-rule walkthrough of tracing the outer loop. We build up the trace in tandem with our walking through of the reduction rules; each snippet that the Record_ rules append to the trace is displayed one by one.

To start, line 1 of the input is matched by Assign, so we reduce by Assign. Line 2 is a **while** loop, which we reduce by While. While converts the loop into an **if** statement testing the condition $x = 0$. This is indeed true by how we mutated the store in line 1, so we can reduce by IfTrue or Trace. In the interest of demonstrating tracing, we reduce by Trace. The trace built thus far is empty, or ϵ . We’ve only entered recording mode, but we haven’t actually recorded any commands yet.

Line 3 in the input is an assignment, which is matched by RecordAssign. RecordAssign appends the assignment itself onto the trace:

Example Trace

```

1  y := 0;

```

Line 4 in the input is the inner loop, and we will now see how the tracing rules deal with recording loops. The loop itself will first reduce to an **if** via RecordWhile, which appends a no-op **skip**; to the trace. In reducing the resulting **if**, we are testing the condition $y = 0$. It is true, so we reduce using RecordIfTrue. The result is that we append a side-exit as a **bail** to the trace. The computation that would have been executed *had the condition been false* gets packaged up as a continuation and gets put into the body of the **bail** (shown indented in the listing):

```

2  skip;
3  bail y ≠ 0 to
4      z := 1;
5      b := a + 1;
6      while x = 0 do
7          y := 0;
8          while y = 0 do
9              y := 1;
10             z := 1;
11             b := a + 1;

```

Now that we have installed the side-exit for entering into the inner loop, we trace the body of the inner loop as straightline code. Line 5 in the input is another assignment, which we record using RecordAssign.

```

12 y := 1;

```

After the body of the inner loop we attempt to reduce the next iteration of that loop. Again, the loop will first reduce to an **if** by RecordWhile. This appends a **skip**; Unlike the last time, however, the condition $y = 0$ is now false, so we instead reduce using RecordIfFalse. We append another side-exit as before, but the packaged continuation is different. Since the condition was false in the actual execution, we need to include the statement that would have been executed if the condition were true. After that statement we package the rest of the iteration of the outer loop and append it:

```

13 skip;
14 bail y = 0 to
15     y := 1;
16     while y = 0 do

```

$$\begin{aligned}
& FV : \text{Statement} \rightarrow \text{Variables} \\
& FV(s) = \{x \mid x \text{ is free in } s\} \\
& F : ((\text{Expression} + \text{Statement} + \text{Command}) \times \text{Store} \times \mathcal{V}) \rightarrow \text{Statement} \\
& F(e, \sigma, v) = \begin{cases} n & \text{if } e = n \\ \sigma(x) \oplus 1 & \text{if } e = x + 1 \wedge x \in v \\ e & \text{if } e = x + 1 \wedge x \notin v \end{cases} \\
& F(s, \sigma, v) = \begin{cases} s & \text{if } s = \epsilon \\ F(c, \sigma, v) F(s_1, \sigma, v) & \text{if } s = c s_1 \end{cases} \\
& F(c, \sigma, v) = \begin{cases} x := F(e, \sigma, v) & \text{if } c = x := e \\ c & \text{otherwise} \end{cases} \\
& O : (\text{Statement} \times \text{Store} \times \mathcal{V}) \rightarrow \text{Statement} \\
& O(s, \sigma) = \begin{cases} \text{while } b \text{ do } F(s_1, \sigma, FV(s_1)) & \text{if } s = \text{while } b \text{ do } s_1 \\ s & \text{otherwise} \end{cases}
\end{aligned}$$

Figure 4. Variable Folding O

```

17     y := 1;
18     z := 1;
19     b := a + 1;
20     while x = 0 do
21       y := 0;
22       while y = 0 do
23         y := 1;
24         z := 1;
25         b := a + 1;

```

Finally, we apply RecordAssign twice to lines 6–7 and append the assignments to the trace.

```

26     z := 1;
27     b := a + 1;

```

Having successfully traced an iteration of the loop, we now reduce by Stitch to stitch the trace back into the program using the identity as the O function. Abbreviating the continuations for the side-exits as k_i , the final stitched traced loop is as follows.

Abbreviated Stitched and Traced Example Loop

```

1  while x = 0 do
2    y := 0;
3    skip;
4    bail y ≠ 0 to k1
5    y := 1;
6    skip;
7    bail y = 0 to k2
8    z := 1;
9    b := a + 1;

```

2.4 Example O

We have seen the output of tracing, but we obviously want to do more than that. We want to optimize. Consider optimization shown in figure 4 that folds away variables that we never assign to inside a traced loop. First we define a function that calculates the “free” (in the sense of never-assigned-to) variables of a statement. It is assumed to be defined in the usual way. Next we define a helper function F that does the actual optimization. \mathcal{V} is the set of free variables. Finally the O function is just a wrapper around F that calculates and passes in the free variables.

If we apply it to our running example where $a \mapsto 41$, we fold the assignment to b on line 9 of the abbreviated stitch example:

```

9     b := 42;

```

The main benefit of run-time optimization is that we can be more aggressive than with ahead-of-time optimization. Here we presented a simple conservative folding of free loop variables. The idea is that free variables in the loop body can be treated as constants and folded until we break out of the loop. We cannot be so bold with a static version of this kind of folding, as we can only do so if we know that the variables we want to fold are constants for the entirety of program execution. Here, however, we only need to know that the variable’s value does not change *until the loop is finished*.⁴

3. Correctness

What does it mean for a trace to be correct? First, correctness of the traced code is behavioral correctness—the trace has to “do the same thing” as the original code. Attempting to prove confluence of the program text such as in Pfenning [16] is unfruitful, as there are no guarantees in trace compilation of the traced code converging back to the same text as the original program. In a reactive user-interface, for instance, we might trace-compile many inner loops, and those compiled inner loops might execute forever, waiting for user input. But even in those cases of infinite execution we still want to reason about correctness. The need for infinite executions suggests the tool of bisimilarity.

Second, correctness of the traced code is intensional correctness. Unlike ahead-of-time compiler correctness, we cannot say that an optimized trace is observationally equivalent, or has the same sequence of observable reductions, to the original loop in the traditional, extensional sense. Specifically, an optimized trace need *not* be observationally equivalent to the original loop under all stores. Consider the following version of our little example:

```

1  x := 0;
2  while x = 0 do
3    b := a + 1;

```

A reasonable trace-based optimization if $a \mapsto 41$, as we have seen, would be to replace the loop body with $b := 42$. But this code is

⁴In our simple model, the trace is effectively discarded after the loop exits. There is no way to re-enter a traced loop once it exits. This is not the case in practice, where constructs such as methods allow compiled traces to be called multiple times. In those cases the tracing JIT has to add in more guards and side-exits to guard the folded values. We omit this complexity.

most definitely not observationally equivalent to the original: the original has a free variable, a , and the optimized code $b := 42$ does not. Side-exits are also problematic. How do we ensure that we jump back to the right place in the original code?

We retain the familiar notion of observational equivalence, but parameterize it over stores and computation suffixes. Namely, a trace is correct if it is observationally equivalent to the original loop *for the store that the original loop is currently reducing under and for the rest of the program that the original loop would have reduced under*.

To formalize these intuitions, we model correctness using intensionalized bisimulations over stores. Intensionalizing to a particular suffix will be made formal in the definition of O soundness in section 3.3. Bisimulation techniques see popular use in process calculi [14]. There is also existing work in the analysis and correctness proofs of program transformation [6, 21, 22].

The definitions here are built upon, but slightly different from, the standard notions found in the concurrency literature [14], as they are defined over a store. Observational equivalence also becomes formally defined as the notion of bisimilarity. Let the set of labels be defined as follows.

$$Act = \{\delta \mid \delta \text{ is a store update}\} \cup \{\tau\}$$

Definition. If $r \in Act^*$, then \hat{r} is the sequence whereby all occurrences of τ are removed.

Definition. If $r = \alpha_1 \cdots \alpha_n \in Act^*$, we write $m \xrightarrow{r} m'$ to mean

$$m \xrightarrow{\tau}^* \cdot \xrightarrow{\alpha_1} \cdot \xrightarrow{\tau}^* \cdots \xrightarrow{\tau}^* \cdot \xrightarrow{\alpha_n} \cdot \xrightarrow{\tau}^* m'$$

That is, there may be any number of intervening silent transitions between the observable sequences. In this particular system, the primary observable entity is the store itself, so the intuitive meaning of a program becomes the sequence of store updates it performs.

We only concern ourselves with closed program-store pairs in this paper, where the definition of *closed* is as follows.

Definition. For a store σ and a program m , we say m is σ -closed if for all variables that appear in m , $\hat{\sigma}(x)$ is defined.

For the rest of the paper, when we say “for any store” or “for all stores”, we mean for all stores that form closed program-store pairs with the programs under consideration.

Definition (Bisimulation). A *bisimulation* for two reduction relations X, Y is a relation \mathcal{R} such that $\mathcal{R}(\sigma, m, n)$ implies

1. Whenever $\langle \sigma, m \rangle \xrightarrow{\alpha} \langle \sigma', m' \rangle$ then, for some n' , $\langle \sigma, n \rangle \xrightarrow{\alpha} \langle \sigma', n' \rangle$ and $\mathcal{R}(\sigma', m', n')$
2. Whenever $\langle \sigma, n \rangle \xrightarrow{\alpha} \langle \sigma', n' \rangle$ then, for some m' , $\langle \sigma, m \rangle \xrightarrow{\alpha} \langle \sigma', m' \rangle$ and $\mathcal{R}(\sigma', m', n')$

In the above definition we abuse notation and let m, m', n , and n' range over both statements and triples of statements. That is, since it does not add to the discussion to distinguish between 2-tuples and 4-tuples in the definition, we use a single metavariable to range over both.

The traditional notion of bisimilarity is a special case of this one: two programs are bisimilar in the traditional sense if they are bisimilar for all stores.

Definition. m is said to be *bisimilar* to n under reduction relations X, Y for a store σ , written $m \approx_Y n$ for σ , if $\mathcal{R}(\sigma, m, n)$ for some bisimulation \mathcal{R} on X, Y . In other words,

$$\approx_Y = \bigcup \{\mathcal{R} \mid \mathcal{R} \text{ is a bisimulation for } X, Y\}$$

Lemma 3.1. Bisimilarity is an equivalence relation.


Proof. Straightforward. \square

Before stating the main lemma, we note that all nondeterministic rules in our system step to the same store. We prove this later in lemma 3.15. For simplicity in stating the main lemma, we simply say that the two branching stores are always the same.


We are now ready to state the main lemma. In the literature, diamond lemmas are usually single diamonds. The trace calculus, however, has a modal flavor with 2-tuples as one mode and 4-tuples as the other mode. As such, our calculus has six diamonds up to symmetry.

Lemma 3.2 (Diamond Lemma). All of the following hold. For diamonds 4–6, $\langle \rho, k_w, t, p \rangle$ is well-formed, a notion we will expound upon in section 3.1.


1. If $R : \langle \rho, p \rangle \xrightarrow{\alpha} \langle \sigma, m \rangle$ and $R' : \langle \rho, p \rangle \xrightarrow{\alpha} \langle \sigma, n \rangle$, then $m \approx_T n$ for σ .




2. If $R : \langle \rho, p \rangle \xrightarrow{\alpha} \langle \sigma, k_w, t, m \rangle$ and $R' : \langle \rho, p \rangle \xrightarrow{\alpha} \langle \sigma, k'_w, t', n \rangle$, then $\langle k_w, t, m \rangle \approx_T \langle k'_w, t', n \rangle$ for σ .




3. If $R : \langle \rho, p \rangle \xrightarrow{\alpha} \langle \sigma, m \rangle$ and $R' : \langle \rho, p \rangle \xrightarrow{\alpha} \langle \sigma, k_w, t, n \rangle$, then $m \approx_T \langle k_w, t, n \rangle$ for σ .




4. If $R : \langle \rho, k_w, t, p \rangle \xrightarrow{\alpha} \langle \sigma, m \rangle$ and $R' : \langle \rho, k_w, t, p \rangle \xrightarrow{\alpha} \langle \sigma, n \rangle$, then $m \approx_T n$ for σ .



5. If $R : \langle \rho, k_w, t, p \rangle \xrightarrow{\alpha} \langle \sigma, k'_w, t', m \rangle$ and $R' : \langle \rho, k_w, t, p \rangle \xrightarrow{\alpha} \langle \sigma, k''_w, t'', n \rangle$, then $\langle k'_w, t', m \rangle \approx_T \langle k''_w, t'', n \rangle$ for σ .



6. If $R : \langle \rho, k_w, t, p \rangle \xrightarrow{\alpha} \langle \sigma, m \rangle$ and $R' : \langle \rho, k_w, t, p \rangle \xrightarrow{\alpha} \langle \sigma, k'_w, t', n \rangle$, then $m \approx_T \langle k'_w, t', n \rangle$ for σ .



This lemma says that should execution branch into two branches, both branches will do the same thing, at least observationally. We aim to use the main lemma to arrive at a more familiar place: confluence of stores, namely corollary 3.18.

The rest of this section is organized as follows. Section 3.1 introduces the idea of well-formedness for the 4-tuples, or the tracing rules. Section 3.2 introduces the correctness criterion of the unoptimized trace. Section 3.3 proves the main lemma. Section 3.4 explores the relationship between confluence and our bisimulation result.

3.1 Well-Formedness of 4-Tuples

When the calculus decides to initiate a trace, it steps to a configuration in the shape of a 4-tuple. The four components are, in order, the store, the point in the original code when we started tracing, the trace so far, and the statement currently being reduced. Not all 4-tuples are created equal, however, as not all 4-tuples are *well-formed*. Intuitively, well-formedness is something like an incremental version of correctness. Only well-formed 4-tuples eventually become fully correct unoptimized traces. Thus, we want it to be an invariant of the computation.

Well-formedness is a tight and intricate relationship between the original loop, the trace thus far, and the current reduct. Informally, we need the trace thus far to be a recording of all the steps that the original loop took just before it reached the current reduct. Before we formally define well-formedness, we formalize what it means to be a “trace thus far”.

Definition (Partial Trace Relation). A *partial trace relation* is a relation \mathcal{T} such that $\mathcal{T}(\sigma, t, l)$ implies that whenever $\langle \sigma, t \rangle \xrightarrow{\alpha}_B \langle \sigma', t' \rangle$ then, for some l' , $\langle \sigma, l \rangle \xrightarrow{\alpha}_B \langle \sigma', l' \rangle$ and

1. If t stepped by `BailTrue`, $t' = l'$
2. Otherwise, $\mathcal{T}(\sigma', t', l')$

The constituents are as follows: t is the trace and l is the original code. Recall that both t and l are just statements. The formalization is a variation on the standard simulation relation and captures the two properties that a partially constructed trace intuitively satisfies. First, `BailTrue` models jumping back to the original code, so we expect the descendants to be exactly equal. Second, the partial trace is *partial*, so it can terminate before the original code does, signifying that the rest has not yet been traced.

Definition. We call t a *partial trace* to l for a store σ , written $t \approx l$ for σ , if $\mathcal{T}(\sigma, t, l)$ for some partial trace relation \mathcal{T} . In other words,

$$\approx = \bigcup \{ \mathcal{T} \mid \mathcal{T} \text{ is a partial trace relation} \}$$

For the definition of well-formedness and subsequent lemmas we will be working with the reduction relation A , which we have not used yet, as well as a notion of being “stuck”. Recall that the reduction relation $A = B \setminus \{\text{BailTrue}\}$.

Definition. For a reduction relation X , we say a configuration $\langle \sigma, m \rangle$ is *not X -stuck* iff $m = \epsilon$ or $\langle \sigma, m \rangle \xrightarrow{\alpha}_X \langle \sigma', m' \rangle$ for some σ', m' .

Definition (Well-Formedness). A 4-tuple $\langle \sigma_0, k_w, t, m \rangle$ is *well-formed* iff all the following hold.

1. $k_w = (\text{while } b_0 \text{ do } l_0) k_0$

2. For all σ , either

$$\langle \sigma, t \rangle \xrightarrow{r}_A^* \langle \sigma', \epsilon \rangle \text{ and } \langle \sigma, l_0 k_w \rangle \xrightarrow{r}_A^* \langle \sigma', m \rangle$$

or for some t' ,

$$\langle \sigma, t \rangle \xrightarrow{r'}_A^* \langle \sigma'', t' \rangle \text{ and } t' \text{ is } A\text{-stuck but not } B\text{-stuck}$$

3. For all σ , $t \approx l_0 k_w$ for σ .

This property formalizes the invariant we wish computation to preserve. First, k_w must be the loop where we started tracing. Second, the trace t must do one of two things. It must either “go far enough” by terminating in the same reduction sequence that body of the original loop undergoes to reduce to the current statement, m , or it must eventually step to some descendant that can only reduce by `BailTrue`. Third, t must be a partial trace to the original loop for all stores.

We now prove that the reduction relation T preserves this invariant. We will do this in two steps. First, we prove that the tracing rules themselves—the rules that step from a 4-tuple to another 4-tuple—preserve well-formedness. Second, we prove that whenever we initiate a trace—whenever we step from a 2-tuple to a 4-tuple—the resulting 4-tuple is well-formed.

Lemma 3.3. $\xrightarrow{*}_B$ is deterministic.

Proof. $\xrightarrow{*}_B$ has no points of nondeterminism. \square

Lemma 3.4. For some m and σ , $\langle \sigma, m \rangle \xrightarrow{\alpha}_A \langle \sigma', m' \rangle$ iff $\langle \sigma, m k \rangle \xrightarrow{\alpha}_A \langle \sigma', m' k \rangle$ for any k .

Proof. By simple case analysis on the reduction relation A . \square

Lemma 3.5. For some t, l, σ where $\langle \sigma, t \rangle \xrightarrow{\alpha}_A \langle \sigma', t' \rangle$ and $\langle \sigma, l \rangle \xrightarrow{\alpha}_A \langle \sigma', l' \rangle$, a partial trace relation \mathcal{T}_1 where $\mathcal{T}_1(\sigma, t, l)$, and another partial trace relation \mathcal{T}_2 where $\mathcal{T}_2(\sigma', t', c, l')$ for all c , there exists a partial trace relation \mathcal{T}_3 such that $\mathcal{T}_3(\sigma, t, c, l)$.

Proof. We wish to exhibit a partial trace relation \mathcal{T}_3 such that $\mathcal{T}_3(\sigma, t, c, l)$. We claim the following relation is a partial trace relation.

$$\mathcal{T}_3 = \{ (\sigma, t, c, l) \mid \mathcal{T}_1(\sigma, t, l) \} \cup \mathcal{T}_2$$

We proceed by using the definition of the partial trace relation. Suppose the left side is $t c$ and the right side is l . The left side steps.

By assumption we have $\langle \sigma, t \rangle \xrightarrow{\alpha}_A \langle \sigma', t' \rangle$. By $\mathcal{T}_1(\sigma, t, l)$ then we know that $\langle \sigma, l \rangle \xrightarrow{\alpha}_B \langle \sigma', l' \rangle$ and either $t' = l'$ if the step was via `BailTrue` or $\mathcal{T}_1(\sigma', t', l')$ if not via `BailTrue`. By the definition of A , we know it was not via `BailTrue`, so $\mathcal{T}_1(\sigma', t', l')$. By lemma 3.4 we also have $\langle \sigma, t c \rangle \xrightarrow{\alpha}_A \langle \sigma', t' c \rangle$. Note that lemma 3.4 goes in both directions, so we can connect the implications in the following fashion.

$$\begin{aligned}
\langle \sigma, t c \rangle \xrightarrow{\alpha}_A \langle \sigma', t' c \rangle &\Rightarrow \langle \sigma, t \rangle \xrightarrow{\alpha}_A \langle \sigma', t' \rangle & (*) \\
&\Rightarrow \langle \sigma, l \rangle \xrightarrow{\alpha}_A \langle \sigma', l' \rangle
\end{aligned}$$

It now remains to show that $\mathcal{T}_3(\sigma', t' c, l')$. But we already have $\mathcal{T}_2(\sigma', t' c, l')$, which we know is a subset of \mathcal{T}_3 by construction, so we are done. \square

Lemma 3.6. For all c and some t, l , and σ where $t \approx l$ for σ , $\langle \sigma, t \rangle \xrightarrow{r}_A^* \langle \sigma', \epsilon \rangle$, $\langle \sigma, l \rangle \xrightarrow{r}_A^* \langle \sigma', l' \rangle$, and $c \approx l'$ for σ' , we have $t c \approx l$ for σ .

Proof. By induction on the multistep relation $\xrightarrow{*}_A$.

Case: \xrightarrow{r}_A^* is the identity.

That is, $t = \epsilon$, $l = l'$, and $\sigma = \sigma'$. Thus, showing $t c \lesssim l$ for σ is equivalent to showing $c \lesssim l'$ for σ' . We have this by assumption, so we are done.

Case: \xrightarrow{r}_A^* is $\xrightarrow{\alpha}_A \cdot \xrightarrow{r'}_A^*$.

By $t \lesssim l$ for σ , we have $\langle \sigma, t \rangle \xrightarrow{\alpha}_A \langle \sigma'', t' \rangle$ implying that $\langle \sigma, l \rangle \xrightarrow{\alpha}_A \langle \sigma'', l'' \rangle$ and $t' \lesssim l''$ for σ'' . We can then apply the induction hypothesis on $\xrightarrow{r'}_A^*$ to get that $t' c \lesssim l''$ for σ'' .

By $t \lesssim l$ for σ we know there exists a partial trace relation \mathcal{T}_1 such that $\mathcal{T}_1(\sigma, t, l)$. By $t' c \lesssim l''$ for σ'' we know there also exists a \mathcal{T}_2 such that $\mathcal{T}_2(\sigma'', t' c, l'')$. Now we can apply lemma 3.5 on \mathcal{T}_1 , \mathcal{T}_2 , and $\xrightarrow{\alpha}_A$ to get that there exists a \mathcal{T}_3 such that $\mathcal{T}_3(\sigma, t c, l)$. But to show that $t c \lesssim l$ for σ it suffices to exhibit a partial trace relation that \mathcal{T} such that $\mathcal{T}(\sigma, t c, l)$, so we are done. \square

Lemma 3.7. For all c and some t, l , and σ where $t \lesssim l$ for σ and

$$\langle \sigma, t \rangle \xrightarrow{r}_A^* \langle \sigma', t' \rangle \xrightarrow{\tau}_B \langle \sigma'', t'' \rangle$$

where the last τ step is BailTrue, $t c \lesssim l$ for σ .

Proof. By induction on the multistep relation \xrightarrow{r}_A^* .

Case: \xrightarrow{r}_A^* is the identity.

That is, $t' = t$, $\sigma' = \sigma$, and $\langle \sigma, t \rangle \xrightarrow{\tau}_B \langle \sigma'', t'' \rangle$. Further by the definition of the partial trace relation we know $\langle \sigma, l \rangle \xrightarrow{\tau}_B \langle \sigma'', l' \rangle$ such that $t'' = l'$. To show that $t c \lesssim l$ for σ it suffices to exhibit a partial trace relation \mathcal{T} such that $\mathcal{T}(\sigma, t c, l)$. We claim the following \mathcal{T} is such a relation.

$$\mathcal{T} = \{(\sigma, t c, l) \mid t \lesssim l \text{ for } \sigma\}$$

Since BailTrue clobbers its continuation, by inversion we know $\langle \sigma, t c \rangle \xrightarrow{\tau}_B \langle \sigma'', t'' \rangle$. We have already seen that $\langle \sigma, l \rangle \xrightarrow{\tau}_B \langle \sigma'', l' \rangle$. Since we stepped via BailTrue, for \mathcal{T} to be a partial trace relation it remains to show $t'' = l'$. This holds, so we are done.

The converse proof is analogous.

Case: \xrightarrow{r}_A^* is $\xrightarrow{\alpha}_A \cdot \xrightarrow{r'}_A^*$.

Analogous to the second case of lemma 3.6. \square

Now we prove that reduction rules that step from a 4-tuple to another 4-tuple preserve well-formedness.

Lemma 3.8. Let $p = \langle \sigma_0, k_w, t, m \rangle$. If p is well-formed and $p \xrightarrow{\alpha}_T p'$ such that p' is a 4-tuple, then p' is also well-formed.

Proof. The first conjunct holds trivially because no rules change k_w .

We proceed to prove the next two conjuncts together by case analysis on the structure of the reduction relation. Let \mathcal{T} be the partial trace relation from the assumption that the premise is well-formed.

Case: RecordAssign. We want to show the following two things.

2. For all σ , either

$$\langle \sigma, t(x := e;) \rangle \xrightarrow{r_2}_A^* \langle \sigma'_2, \epsilon \rangle \text{ and } \langle \sigma, l_0 k_w \rangle \xrightarrow{r_2}_A^* \langle \sigma'_2, k \rangle$$

or for some t'_2 ,

$$\langle \sigma, t(x := e;) \rangle \xrightarrow{r'_2}_A^* \langle \sigma''_2, t'_2 \rangle$$

where t' is A -stuck but not B -stuck.

3. For all σ , $t(x := e;) \lesssim l_0 k_w$ for σ .

We proceed by case analysis on conjunct 2 of the premise. We have the following subcases for all σ .

Subcase: $\langle \sigma, t \rangle \xrightarrow{r_1}_A^* \langle \sigma'_1, \epsilon \rangle$ and $\langle \sigma, l_0 k_w \rangle \xrightarrow{r_1}_A^* \langle \sigma'_1, x := e; k \rangle$.

2. By lemma 3.4 and applying Assign we know that

$$\langle \sigma, t(x := e;) \rangle \xrightarrow{r_1}_A^* \langle \sigma'_1, x := e; \rangle \xrightarrow{\delta}_A \langle \sigma'_1[x/\hat{\sigma}'_1(e)], \epsilon \rangle$$

Similarly we can apply Assign and get

$$\langle \sigma, l_0 k_w \rangle \xrightarrow{r_1}_A^* \langle \sigma'_1, x := e; k \rangle \xrightarrow{\delta}_A \langle \sigma'_1[x/\hat{\sigma}'_1(e)], k \rangle$$

3. We first show that $x := e; \lesssim x := e; k$ for σ'_1 and then use lemma 3.6. It suffices to demonstrate a partial trace relation \mathcal{T} such that $\mathcal{T}(\sigma'_1, x := e; , x := e; k)$. We claim the following is one such relation.

$$\mathcal{T} = \{(\sigma'_1, x := e; , x := e; k) \cup \{(\sigma'_1[x/\hat{\sigma}'_1(e)], \epsilon, k)\}$$

Suppose the left side is $x := e$. We can apply Assign to the left side to get:

$$\langle \sigma'_1, x := e; \rangle \xrightarrow{\delta}_A \langle \sigma'_1[x/\hat{\sigma}'_1(e)], \epsilon \rangle \text{ where } \delta = x/\hat{\sigma}'_1(e)$$

Obviously we can also apply BailTrue to the right side to get

$$\langle \sigma'_1, x := e; k \rangle \xrightarrow{\delta}_A \langle \sigma'_1[x/\hat{\sigma}'_1(e)], k \rangle \text{ where } \delta = x/\hat{\sigma}'_1(e)$$

Now it only remains to show $\mathcal{T}(\sigma'_1[x/\hat{\sigma}'_1(e)], \epsilon, k)$ holds. But this is already in \mathcal{T} by construction.

Suppose the left side is ϵ , it does not step, so we are done.

We apply lemma 3.6 on $x := e; \lesssim x := e; k$ for σ'_1 and $\xrightarrow{r_1}_A^*$ to get the desired result.

Subcase: For some t'_1 , $\langle \sigma, t \rangle \xrightarrow{r'_1}_A^* \langle \sigma''_1, t'_1 \rangle$ and t'_1 is A -stuck but not B -stuck.

2. By lemma 3.4 we know that

$$\langle \sigma, t(x := e;) \rangle \xrightarrow{r'_1}_A^* \langle \sigma''_1, t'_1(x := e;) \rangle$$

We know t'_1 is A -stuck but not B -stuck, so by inversion on each reduction rule $t'_1(x := e;)$ is still A -stuck but not B -stuck.

3. By inversion, if t'_1 is A -stuck but not B -stuck, the only rule it can reduce by is BailTrue. So, we have

$$\langle \sigma, t \rangle \xrightarrow{r'_1}_A^* \langle \sigma''_1, t'_1 \rangle \xrightarrow{\tau}_B \langle \sigma'''_1, t'_1 \rangle$$

We apply lemma 3.7 to get the desired result.

Case: RecordSeq. This case is analogous to the case of RecordAssign.

Case: RecordIfFalse. We want to show the following two things.

2. For all σ , either

$$\langle \sigma, t(\mathbf{bail} \ b \ \mathbf{to} \ (s \ k)) \rangle \xrightarrow{r_2}_A^* \langle \sigma'_2, \epsilon \rangle \text{ and } \langle \sigma, l_0 k_w \rangle \xrightarrow{r_2}_A^* \langle \sigma'_2, k \rangle$$

or for some t' ,

$$\langle \sigma, t(\mathbf{bail} \ b \ \mathbf{to} \ (s \ k)) \rangle \xrightarrow{r'_2}_A^* \langle \sigma''_2, t'_2 \rangle$$

where t' is A -stuck but not B -stuck.

3. For all σ , $t(\mathbf{bail} \ b \ \mathbf{to} \ (s \ k)) \lesssim l_0 k_w$ for σ .

We proceed by case analysis on conjunct 2 of the premise. We have the following subcases for all σ .

Subcase: $\langle \sigma, t \rangle \xrightarrow{r}_A^* \langle \sigma'_1, \epsilon \rangle$ and $\langle \sigma, l_0 k_w \rangle \xrightarrow{r_1}_A^* \langle \sigma'_1, (\mathbf{if } b \mathbf{ then } s) \rangle k$.

2. By lemma 3.4 we know that

$$\langle \sigma, t (\mathbf{bail } b \mathbf{ to } (s k)) \rangle \xrightarrow{r_1}_A^* \langle \sigma'_1, \mathbf{bail } b \mathbf{ to } (s k) \rangle$$

Suppose at this point we can apply BailTrue, then we have exhibited a descendant such that it is A -stuck but not B -stuck. If instead we can apply BailFalse, then we have

$$\langle \sigma'_1, \mathbf{bail } b \mathbf{ to } (s k) \rangle \xrightarrow{\tau}_A \langle \sigma'_1, \epsilon \rangle$$

In this case by inversion we know $\hat{\sigma}'_1(b) = \mathit{false}$, so we can apply IfFalse to get

$$\begin{aligned} \langle \sigma, l_0 k_w \rangle &\xrightarrow{r_1}_A^* \langle \sigma'_1, (\mathbf{if } b \mathbf{ then } s) \rangle k \\ &\xrightarrow{\tau}_A \langle \sigma'_1, k \rangle \end{aligned}$$

3. We first show that $\mathbf{bail } b \mathbf{ to } (s k) \approx (\mathbf{if } b \mathbf{ then } s) k$ for σ'_1 and then use lemma 3.6. It suffices to demonstrate a partial trace relation \mathcal{T} such that $\mathcal{T}(\sigma'_1, \mathbf{bail } b \mathbf{ to } (s k), (\mathbf{if } b \mathbf{ then } s) k)$. We claim the following is one such relation.

$$\begin{aligned} \mathcal{T} = &\{ \langle \sigma'_1, \mathbf{bail } b \mathbf{ to } (s k), (\mathbf{if } b \mathbf{ then } s) k \rangle \} \\ &\cup \{ \langle \sigma'_1, \epsilon, k \rangle \} \end{aligned}$$

Suppose the left side is **bail** and that we can apply BailFalse.

$$\langle \sigma'_1, \mathbf{bail } b \mathbf{ to } (s k) \rangle \xrightarrow{\tau}_B \langle \sigma'_1, \epsilon \rangle$$

Then have by inversion that $\hat{\sigma}'_1(b) = \mathit{false}$, so we can apply IfFalse to the right side:

$$\langle \sigma'_1, (\mathbf{if } b \mathbf{ then } s) k \rangle \xrightarrow{\tau}_B \langle \sigma'_1, k \rangle$$

It remains to show $\mathcal{T}(\sigma'_1, \epsilon, k)$ holds. But this is already in \mathcal{T} by construction.

Suppose the left side is **bail** and that we can apply BailTrue.

$$\langle \sigma'_1, \mathbf{bail } b \mathbf{ to } (s k) \rangle \xrightarrow{\tau}_B \langle \sigma'_1, s k \rangle$$

Then we have by inversion that $\hat{\sigma}'_1(b) = \mathit{true}$, so we can apply IfTrue to the right side:

$$\langle \sigma'_1, (\mathbf{if } b \mathbf{ then } s) k \rangle \xrightarrow{\tau}_B \langle \sigma'_1, s k \rangle$$

Since the left side stepped via BailTrue and $s k = s k$, we are still in \mathcal{T} .

Suppose the left side is ϵ . It does not step, so we are done.

We apply lemma 3.6 on $\mathbf{bail } b \mathbf{ to } (s k) \approx (\mathbf{if } b \mathbf{ then } s) k$ for σ'_1 and $\xrightarrow{r_1}_A^*$ to get the desired result.

Subcase: For some t'_1 , $\langle \sigma, t \rangle \xrightarrow{r'_1}_A^* \langle \sigma', t'_1 \rangle$ and t'_1 is A -stuck but not B -stuck.

Analogous to the argument presented in the same subcase for RecordAssign.

Case: RecordIfTrue. This case is analogous to the case of RecordIfFalse, except with boolean conditions reversed.

Case: RecordWhile. The proof proceeds analogously to the case of RecordAssign and RecordSeq except for the following. For the left side of t (**skip**;) we step using Seq and for the right side we step using While. \square

Now it remains to prove that whenever we initiate a trace, the resulting 4-tuple is well-formed.

Lemma 3.9. If $\langle \sigma, m \rangle \xrightarrow{\alpha}_T p'$ where p' is a 4-tuple, then p' is well-formed.

Proof. By inversion the only rule that results in a 4-tuple is Trace. Let $w = \mathbf{while } b \mathbf{ do } s$. We have

$$\langle \sigma, (\mathbf{if } b \mathbf{ then } (s w)) k \rangle \xrightarrow{\tau}_T \langle \sigma, w k, \epsilon, s w k \rangle$$

The first two conjuncts are clearly satisfied. It remains to prove conjunct 3, that $\epsilon \approx s w k$ for all σ . It suffices to exhibit a partial trace relation \mathcal{T} such that for all σ , $\mathcal{T}(\sigma, \epsilon, s w k)$. Since $t = \epsilon$, we exhibit the empty relation \emptyset as one such \mathcal{T} . \square

Lemmas 3.8 and 3.9 lead us to a more general lemma about the transitive, reflexive closure of the T reduction relation. This lemma is not used in the rest of the section, but does clearly convey that well-formedness is a property preserved by computation in our calculus.

Lemma 3.10. If $\langle \sigma, m \rangle \xrightarrow{r}_T^* p'$ where p' is a 4-tuple, then p' is well-formed.

Proof. By induction on structure of \xrightarrow{r}_T^* .

Case: \xrightarrow{r}_T^* is the identity. This case is impossible. No 4-tuple is the identity of a 2-tuple.

Case: \xrightarrow{r}_T^* is $\xrightarrow{r'}_T^* \cdot \xrightarrow{r''}_T$. Here we have four subcases.

Subcase: $\xrightarrow{r''}_T$ steps from a 2-tuple to a 2-tuple. This case is impossible, as we have by assumption that p' is a 4-tuple.

Subcase: $\xrightarrow{r''}_T$ steps from a 2-tuple to a 4-tuple. We use lemma 3.9 to obtain the desired result.

Subcase: $\xrightarrow{r''}_T$ steps from a 4-tuple to a 2-tuple. This case is impossible, as we have by assumption that p' is a 4-tuple.

Subcase: $\xrightarrow{r''}_T$ steps from a (left-hand) 4-tuple to a 4-tuple. We first apply the induction hypothesis to obtain that the left-hand 4-tuple is well-formed. We then use lemma 3.8 on the left-hand 4-tuple to obtain the desired result. \square

3.2 Correctness of the Unoptimized Trace

Recall that the intuition for well-formedness is that it is an incremental correctness. With it we can now build up a bisimulation relation.⁵

The following lemma will be useful.

Lemma 3.11. For some m , $\langle \sigma, m \rangle \xrightarrow{*}_B \langle \sigma', m' \rangle$ such that $\langle \sigma', m' \rangle$ does not step iff $m' = \epsilon$.

Proof. By induction on the structure of the reduction relation. \square

Lemma 3.12 (Stitch Lemma). For some t , let $w = \mathbf{while } b_0 \mathbf{ do } l_0$ and $w' = \mathbf{while } b_0 \mathbf{ do } t$. If for some k , $t \approx l_0 w k$ for all σ and (*) holds of $t, l_0 w k, \sigma$ then, $w' k \approx_B w k$ for all σ .

(*) For all σ , either

$$\langle \sigma, t \rangle \xrightarrow{r}_A^* \langle \sigma', \epsilon \rangle \text{ and } \langle \sigma, l_0 w k \rangle \xrightarrow{r}_A^* \langle \sigma', w k \rangle$$

or for some t' ,

$$\langle \sigma, t \rangle \xrightarrow{r'}_A^* \langle \sigma'', t' \rangle \text{ and } t' \text{ is } A\text{-stuck but not } B\text{-stuck}$$

⁵The unoptimized trace is in fact strongly bisimilar to the original code. Since we are simply recording some execution path command-for-command, it shouldn't be surprising that the resulting trace is exactly equivalent to the original path. In the interest of less mechanism and since weak bisimilarity subsumes strong bisimilarity, we will directly prove weak bisimilarity.

This lemma is the correctness property we want to express of unoptimized traces. In prose, w is the original loop, and w' is the new loop with the trace stitched in. The $t \approx l_0 w k$ for all σ and (*) conditions are what hold of $t, l_0 w k, \sigma$ per well-formedness at the point of stitching.

The lemma says once we come full circle and stitch the recorded trace into the original program, that trace is actually equivalent to the original loop. The reader should bear in mind that this is an almost extensional equivalence, a stronger notion than intensional equivalence. The insight here is since we have proven a more restrictive property than we need, we can relax it. For fruitful optimization we need to make the relation larger, relaxing extensional bisimilarity to the intensional version.

Proof. To show that $w' k \approx_B w k$ for all σ , it suffices to exhibit a bisimulation relation \mathcal{R} under the relations B, B such that $\mathcal{R}(\sigma, w' k, w k)$ for all σ . We claim the following relation, for any s , is a bisimulation relation for all stores.

$$\begin{aligned} \mathcal{R} = & \{(\sigma, s, s)\} \\ & \cup \{(\sigma, w' k, w k)\} \\ & \cup \{(\sigma, \text{if } b_0 \text{ then } (t w') k, \text{if } b_0 \text{ then } (l_0 w) k)\} \\ & \cup \{(\sigma, m w' k, n) \mid m \approx n \text{ for } \sigma \text{ and } (*) \text{ holds of } m, n, \sigma\} \end{aligned}$$

We proceed by case analysis on the left-side reduction step.

Case: The left side and right side are the same.

By lemma 3.3 B is deterministic, both sides step using the same rule, producing the same descendants. But then they are in \mathcal{R} by construction.

$$\begin{array}{ccc} \langle \sigma, s \rangle & \mathcal{R} & \langle \sigma, s \rangle \\ \downarrow \alpha & & \downarrow \alpha \\ \langle \sigma', s' \rangle & \mathcal{R} & \langle \sigma', s' \rangle \end{array}$$

Case: The left side is $w' k$ and the right side is $w k$.

Under all stores, both $w' k$ and $w k$ step using **While**. The resulting two statements of the above steps are in \mathcal{R} by construction.

$$\begin{array}{ccc} \langle \sigma, w' k \rangle & \mathcal{R} & \langle \sigma, w k \rangle \\ \downarrow \text{While } \tau & & \downarrow \tau \text{ While} \\ \langle \sigma, \text{if } b_0 \text{ then } (t w') k \rangle & \mathcal{R} & \langle \sigma, \text{if } b_0 \text{ then } (l_0 w) k \rangle \end{array}$$

Case: The left side is **if** b_0 **then** $(t w') k$ and the right side is **if** b_0 **then** $(l_0 w) k$. The left-hand side steps using **lffalse**.

We can fill the diagram by inversion as follows.

$$\begin{array}{ccc} \langle \sigma, \text{if } b_0 \text{ then } (t w') k \rangle & \mathcal{R} & \langle \sigma, \text{if } b_0 \text{ then } (l_0 w) k \rangle \\ \downarrow \text{lffalse } \tau & & \downarrow \tau \text{ lffalse} \\ \langle \sigma, k \rangle & \mathcal{R} & \langle \sigma, k \rangle \end{array}$$

Case: The left side is **if** b_0 **then** $(t w') k$ and the right side is **if** b_0 **then** $(l_0 w) k$. The left side steps side using **lffalse**.

We want to show that $\mathcal{R}(\sigma, t w' k, l_0 w k)$. We already have by assumption that $t \approx l_0 w k$ for σ and (*) holds of $t, l_0 w k, \sigma$, so $\mathcal{R}(\sigma, t w' k, l_0 w k)$ holds by construction.

$$\begin{array}{ccc} \langle \sigma, \text{if } b_0 \text{ then } (t w') k \rangle & \mathcal{R} & \langle \sigma, \text{if } b_0 \text{ then } (l_0 w) k \rangle \\ \downarrow \text{lffalse } \tau & & \downarrow \tau \text{ lffalse} \\ \langle \sigma, t w' k \rangle & \mathcal{R} & \langle \sigma, l_0 w k \rangle \end{array}$$

Case: The left side is $m w' k$ and the right side is n such that $m \approx n$ for σ and (*) holds of m, n, σ . The left side steps using **BailTrue**.

Clearly, since $m w' k$ steps using **BailTrue**, m steps using **BailTrue**. By conjunct 1a the definition of the partial trace relation we can fill out the diagram as below. Since the descendants are equal by conjunct 1a, they are still in \mathcal{R} .

$$\begin{array}{ccc} \langle \sigma, m w' k \rangle & \mathcal{R} & \langle \sigma, n \rangle \\ \downarrow \text{BailTrue } \tau & & \downarrow \tau \\ \langle \sigma, n' \rangle & \mathcal{R} & \langle \sigma, n' \rangle \end{array}$$

Case: The left side is $m w' k$ and the right side is n such that $m \approx n$ for σ and (*) holds of m, n, σ . The left side steps using a rule other than **BailTrue**.

Clearly, since $m w' k$ steps not using **BailTrue**, m steps without clobbering $w' k$. By conjunct 1b the definition of the partial trace relation we can fill out the diagram as below. We know $m' \approx n'$ for σ' holds by conjunct 1b. By lemma 3.3 we also know that since m' and n' are descendants of m and n , (*) also holds of m', n', σ' . But then $\mathcal{R}(\sigma', m' w' k, n')$ holds by construction.

$$\begin{array}{ccc} \langle \sigma, m w' k \rangle & \mathcal{R} & \langle \sigma, n \rangle \\ \downarrow \alpha & & \downarrow \alpha \\ \langle \sigma', m' w' k \rangle & \mathcal{R} & \langle \sigma', n' \rangle \end{array}$$

The proofs for the converses are symmetric except for the following two cases.

Case: The left side is $m w' k$ and the right side is n such that $m \approx n$ for σ and (*) holds of m, n, σ . The right side n takes a step and the left side m takes a step.

We know from the prefix trace relation that if m steps, then n takes the same step. Together with lemma 3.3, this implies that should both side step, then both sides take the exact same step, so we are done.

Case: The left side is $m w' k$ and the right side is n such that $m \approx n$ for σ and (*) holds of m, n, σ . The right side n takes a step but the left side m does not.

Since ϵ does not step, by (*) we conclude that \xrightarrow{r}_A^* is the identity. But then we also know that $n = w k$. It remains to show that $\mathcal{R}(\sigma, w' k, w k)$, which holds by construction. \square

3.3 Proof of the Diamond Lemma

Per usual, we start with some helper lemmas.

Lemma 3.13 (Inclusion Lemma). $\longrightarrow_B \subset \longrightarrow_T$

Proof. Trivial. \square

Lemma 3.14. If $\langle \sigma, k_w, t, m \rangle \xrightarrow{\alpha}_T \langle \sigma', k_w, t', m' \rangle$ then, $\langle \sigma, m \rangle \xrightarrow{\alpha}_B \langle \sigma', m' \rangle$.

Proof. By straightforward case analysis on the structure of $\xrightarrow{\alpha}_T$. \square


With the tool of bisimilarity under our belt, we can now make precise the notion of the soundness of O . This definition will be crucial for the proof of the diamond lemma.

Definition (O -Soundness). An O function is *sound* iff for any w, w', k, σ such that $w' k \approx_B w k$ for all stores, $O(w', \sigma) k \approx_B w k$ for σ .

Proof of the diamond lemma. If R and R' are the same, then we are done as $m = n$ or $\langle k_w, t, n \rangle = \langle k'_w, t', n' \rangle$. It is straightforward to verify that $R = R'$ for diamonds 1, 2, 4, and 5, which are deterministic, so we only need to prove diamonds 3 and 6. We can rewrite these diamonds more precisely below.


3. The nondeterministic rules are `IfTrue` and `Trace`. We have $m = n$ by inversion. Further, by lemma 3.9 $\langle \sigma, k_w, t, n \rangle$ is well-formed.

If $R : \langle \rho, p \rangle \xrightarrow{\alpha}_T \langle \sigma, m \rangle$ and $R' : \langle \rho, p \rangle \xrightarrow{\alpha}_T \langle \sigma, k_w, t, n \rangle$, then $m \approx_T \langle k_w, t, n \rangle$ for σ .



6. The nondeterministic rules are `Record_` and `Abort`. We have $k_w = k'_w$ and $m = n$ by inversion. Similarly, by lemma 3.8 $\langle \sigma, k_w, t', n \rangle$ is well-formed.

If $R : \langle \rho, w, t, p \rangle \xrightarrow{\alpha}_T \langle \sigma, m \rangle$ and $R' : \langle \rho, k_w, t, p \rangle \xrightarrow{\alpha}_T \langle \sigma, k'_w, t', n \rangle$, then $m \approx_T \langle k'_w, t', n \rangle$ for σ .



To show that \approx_T holds for both diamonds, it suffices to exhibit a bisimulation relation \mathcal{R} under the reductions T, T such that $\mathcal{R}(\sigma, m, \langle k_w, t, n \rangle)$ and $\mathcal{R}(\sigma, m, \langle k_w, t', n \rangle)$ hold.

We claim the following relation is a bisimulation for any $m, n, u, k_v, k_w, t, \sigma$.


$$\begin{aligned}
 \mathcal{R} = & \{ \langle \sigma, m, n \rangle \mid m \approx_B n \text{ for } \sigma \} \\
 & \cup \{ \langle \sigma, m, \langle k_v, u, n \rangle \rangle \mid m \approx_B n \text{ for } \sigma \text{ and } \\
 & \quad \langle \sigma, k_v, u, n \rangle \text{ is well-formed} \} \\
 & \cup \{ \langle \sigma, \langle k_w, t, m \rangle, n \rangle \mid m \approx_B n \text{ for } \sigma \text{ and } \\
 & \quad \langle \sigma, k_w, t, m \rangle \text{ is well-formed} \}
 \end{aligned}$$

Recall that by reflexivity of bisimilarity, $m = n$ implies $m \approx_B n$.

We proceed by case analysis on the left-side reduction step.


Case: The left side and right side are both 2-tuples, and the left reduction step is from a 2-tuple to a 2-tuple.

Since we know $m \approx_B n$ for σ by assumption, we know there exists some n' such that $\langle \sigma, n \rangle \xrightarrow{\hat{\alpha}}_B \langle \sigma', n' \rangle$ such that $m' \approx_B n'$ for σ' . By the inclusion lemma, we have $\langle \sigma, n \rangle \xrightarrow{\hat{\alpha}}_T \langle \sigma', n' \rangle$. Since $m' \approx_B n'$ for σ' , $\mathcal{R}(\sigma', m', n')$ holds.



Case: The left side and the right side are both 2-tuples, and the left reduction step is from a 2-tuple to a 4-tuple.

By inversion the left reduction must be using `Trace` with label τ . By lemma 3.9 `Trace` produces a well-formed 4-tuple. Also by inversion, we know that $S : \langle \sigma, m \rangle \xrightarrow{\tau}_B \langle \sigma, k_w, t, m' \rangle$ using `IfTrue`. Using the assumption $m \approx_B n$ for σ on S and the inclusion lemma, we have $\langle \sigma, n \rangle \xrightarrow{\hat{\tau}}_T \langle \sigma', n' \rangle$ such that $m' \approx_B n'$ for σ' . But then $\mathcal{R}(\sigma', \langle k_w, t, m' \rangle, n')$ holds by construction.




Case: The left side is a 2-tuple, and the right side is a 4-tuple. The left reduction step is from a 2-tuple to a 2-tuple.

Using the assumption $m \approx_B n$ for σ , we have $\langle \sigma, n \rangle \xrightarrow{\hat{\alpha}}_B \langle \sigma', n' \rangle$ such that $m' \approx_T n'$ for σ' .

If the weak step decomposes to $\langle \sigma, n \rangle \xrightarrow{\tau}_B \cdot \xrightarrow{\tau}_B^* \cdot \xrightarrow{\alpha}_B \cdot \xrightarrow{\tau}_B^* \langle \sigma', n' \rangle$, we can use `Abort` to initially step $\langle \sigma, k_v, u, n \rangle$ using the initial τ step as the assumption in the decomposition. The rest then of the reduction then is obtained via the inclusion lemma, giving us $\langle \sigma, k_v, u, n \rangle \xrightarrow{\hat{\alpha}}_T \langle \sigma', n' \rangle$

If the weak step decomposes to $\langle \sigma, n \rangle \xrightarrow{\alpha}_B \cdot \xrightarrow{\tau}_B^* \langle \sigma', n' \rangle$, we can obtain $\langle \sigma, k_v, u, n \rangle \xrightarrow{\hat{\alpha}}_T \langle \sigma', n' \rangle$ analogously by applying `Abort` to the α step.

Since $m' \approx_T n'$ for σ' , $\mathcal{R}(\sigma', m', n')$ holds.



Case: The left side is a 2-tuple, and the right side is a 4-tuple. The left reduction step is from a 2-tuple to a 4-tuple.

By inversion the left reduction must be using `Trace` with label τ . By lemma 3.9 `Trace` produces a well-formed 4-tuple. Also by inversion, we know that $S : \langle \sigma, m \rangle \xrightarrow{\tau}_B \langle \sigma, k_w, t, m' \rangle$ using `IfTrue`. Using the assumption $m \approx_B n$ for σ on S and the technique from the previous case, we have $\langle \sigma, k_v, u, n \rangle \xrightarrow{\hat{\tau}}_T \langle \sigma', n' \rangle$ such that $m' \approx_B n'$ for σ . But then $\mathcal{R}(\sigma, \langle k_w, t, m' \rangle, n')$ holds by constructions.

$$\begin{array}{ccc}
\langle \sigma, m \rangle & \mathcal{R} & \langle \sigma, k_v, u, n \rangle \\
\text{Trace } \tau \downarrow & & \hat{\tau} \downarrow \text{Abort} \\
\langle \sigma, k_w, t, m' \rangle & \mathcal{R} & \langle \sigma, n' \rangle
\end{array}$$

Case: The left side is a 4-tuple, and the right side is a 2-tuple. The left reduction step is from a 4-tuple to a 2-tuple by way of Abort.

The premise of the Abort rule gives us $\langle \sigma, m \rangle \xrightarrow{\alpha}_B \langle \sigma', m' \rangle$. All 2-tuple to 2-tuple reduction rules in T are also in B , so we also have $S : \langle \sigma, m \rangle \xrightarrow{\alpha}_B \langle \sigma', m' \rangle$. Using the assumption $m \approx_B n$ for σ on S and the inclusion lemma, we have $\langle \sigma, n \rangle \xrightarrow{\hat{\alpha}}_T \langle \sigma', n' \rangle$ such that $m' \approx_B n'$ for σ' . $\mathcal{R}(\sigma', m', n')$ then holds by construction.

$$\begin{array}{ccc}
\langle \sigma, k_w, t, m \rangle & \mathcal{R} & \langle \sigma, n \rangle \\
\text{Abort } \alpha \downarrow & & \hat{\alpha} \downarrow \\
\langle \sigma', m' \rangle & \mathcal{R} & \langle \sigma', n' \rangle
\end{array}$$

Case: The left side is a 4-tuple, and the right side is a 2-tuple. The left reduction step is from a 4-tuple to a 2-tuple by way of Stitch.

We have by assumption that $\langle \sigma, k_w, t, m \rangle$ is well-formed. By inversion we know that $k_w = m = (\mathbf{while} \ b \ \mathbf{do} \ s) \ k$, $m' = O(\mathbf{while} \ b \ \mathbf{do} \ t, \sigma) \ k$. Let the $w = \mathbf{while} \ b \ \mathbf{do} \ s$ and $w' = \mathbf{while} \ b \ \mathbf{do} \ t$. By well-formedness we know $t \approx s \ w \ k$ for all σ and lemma 3.12's (*) condition holds of $t, s \ w \ k, \sigma$. We can thus apply lemma 3.12 on $w, w', t \approx s \ w \ k$ for all σ , and (*) to get that $w' \ k \approx_B w \ k$ for all σ . Applying this to the soundness of O , gives $O(w', \sigma) \ k \approx_B w \ k$ for σ , or $m' \approx_B m$. We know by assumption that $m \approx_B n$ for σ , so by transitivity we have $m' \approx_B n$ for σ . We then use Id to complete the diagram. Since $\mathcal{R}(\sigma, m', n)$ holds, we are done.

$$\begin{array}{ccc}
\langle \sigma, k_w, t, (\mathbf{while} \ b \ \mathbf{do} \ s) \ k \rangle & \mathcal{R} & \langle \sigma, n \rangle \\
\text{Stitch } \tau \downarrow & & \hat{\tau} \downarrow Id \\
\langle \sigma, O(\mathbf{while} \ b \ \mathbf{do} \ t, \sigma) \ k \rangle & \mathcal{R} & \langle \sigma, n \rangle
\end{array}$$

Case: The left side is a 4-tuple, and the right side is a 2-tuple. The left reduction step is from a 4-tuple to a 4-tuple by way of a Record_rule.

Using the assumption $m \approx_B n$ for σ , lemma 3.14, and the inclusion lemma, we have $\langle \sigma, n \rangle \xrightarrow{\hat{\alpha}}_T \langle \sigma', n' \rangle$ such that $m' \approx_B n'$ for σ' . By lemma 3.8 $\langle \sigma', k_w, t', m' \rangle$ is well-formed. But then $\mathcal{R}(\sigma', \langle k_w, t', m' \rangle, n')$ holds by construction.

$$\begin{array}{ccc}
\langle \sigma, k_w, t, m \rangle & \mathcal{R} & \langle \sigma, n \rangle \\
\alpha \downarrow & & \hat{\alpha} \downarrow \\
\langle \sigma', k_w, t', m' \rangle & \mathcal{R} & \langle \sigma', n' \rangle
\end{array}$$

The proofs for the converses are symmetric.

This demonstrates that \mathcal{R} is a bisimulation for the programs found in all diamonds, and we are done. \square

3.4 From Bisimulation to Confluence

This section aims to be the interface between bisimulation and confluence. As correctness is often studied in terms of determinacy and confluence, we seek here to prove something akin to confluence of stores to show the adequacy of our operational semantics. Traditionally, confluence theorems are proven from the bottom up using a diamond lemma, iterating that diamond lemma to build a strip lemma, and finally using the strip lemma to construct confluence [16]. Bisimilarity, however, allows us to skip the iteration of the single-step diamond lemma. Indeed, there is no analog to an iterable diamond lemma here. We instead use bisimilarity to directly obtain a strip lemma. Nevertheless, the techniques and diagrams in this section are strongly influenced by the clear and readable approach of Pfenning [16].

First we prove the assumption needed for all cases of the diamond lemma, that nondeterministic branching always branches to configurations with the same store.

Lemma 3.15. If $\langle \sigma, m \rangle \xrightarrow{\alpha}_T \langle \sigma', m' \rangle$ and $\langle \sigma, m \rangle \xrightarrow{\alpha'}_T \langle \sigma'', m'' \rangle$, then $\sigma' = \sigma''$ and $\alpha = \alpha'$.

Proof. Straightforward case analysis. \square

Lemma 3.16 (Strip Lemma). If $R : \langle \sigma, m \rangle \rightarrow_T \langle \sigma', m' \rangle$ and $R^{*'} : \langle \sigma, m \rangle \rightarrow_{T'}^* \langle \sigma'', m'' \rangle$, then for some ρ, n', n'' , $\langle \sigma', m' \rangle \rightarrow_{T'}^* \langle \rho, n' \rangle$ and $\langle \sigma'', m'' \rangle \rightarrow_T^* \langle \rho, n'' \rangle$ such that $n' \approx_T n''$ for ρ .

$$\begin{array}{ccc}
& \langle \sigma, m \rangle & \\
R \swarrow & & \searrow R^{*'} \\
\langle \sigma', m' \rangle & & \langle \sigma'', m'' \rangle \\
S^* \downarrow & & \downarrow S^{*'} \\
\langle \rho, n' \rangle & \approx_T & \langle \rho, n'' \rangle
\end{array}$$

Proof. By case analysis on the structure of R^* and the definition of \approx_T .

Case: R^* is the identity reduction.


Then $\langle \sigma'', m'' \rangle = \langle \sigma, m \rangle$, and we can let $\rho = \sigma'$ and fill out the diagram as follows.

$$\begin{array}{ccc}
& \langle \sigma, m \rangle & \\
R \swarrow & & \searrow Id \\
\langle \sigma', m' \rangle & & \langle \sigma, m \rangle \\
Id \downarrow & & \downarrow R \\
\langle \sigma', m' \rangle & \approx_T & \langle \sigma', m' \rangle
\end{array}$$


Case: $R^{*'}$ ends in a reduction step R_1' followed by $R_2^{*'}$.

From lemma 3.9 we know m_1'' is well-formed if it is a 4-tuple. We can appeal to lemmas 3.15 and then the diamond lemma on R_1' and R to get that $\langle \sigma', m' \rangle \approx_T \langle \sigma', m_1'' \rangle$.⁶ Since we have by assumption $\langle \sigma', m_1'' \rangle \xrightarrow{*}_T \langle \sigma'', m'' \rangle$, we can iteratively use the first condition of the definition of bisimulation to obtain an S^* .

⁶Enumerating six diamonds detracts from readability while adding no explanatory aid, so here we abuse notation and only list one diamond.




□



Now we can prove the diamond property for stores on the multistep reduction, which we call *store confluence*. □


Theorem 3.17 (Store Confluence). If $R^* : \langle \sigma, m \rangle \rightarrow_T^* \langle \sigma', m' \rangle$ and $R^{*'} : \langle \sigma, m \rangle \rightarrow_T^* \langle \sigma'', m'' \rangle$, then for some ρ, n', n'' , $\langle \sigma', m' \rangle \rightarrow_T^* \langle \rho, n' \rangle$ and $\langle \sigma'', m'' \rangle \rightarrow_T^* \langle \rho, n'' \rangle$ such that $n' \approx_T n''$ for ρ .



Proof. By induction on the structure of R^* .

Case: R^* is the identity reduction.

Then $\langle \sigma', m' \rangle = \langle \sigma, m \rangle$ and we can fill out the diagram as follows.



Case: R^* ends in a reduction step R_1 followed by R_2^* .

We can apply the strip lemma on R_1 and $R^{*'}$ to get $S_1^{*''}$ and $S_1^{*'}$. Then by the induction hypothesis on R_2^* and $S_1^{*''}$ we get S^* and $S_2^{*''}$. Now that we have $S_2^{*''}$, we can iteratively apply the first condition of the definition of bisimulation to obtain an $S_2^{*'}$. By transitivity of bisimilarity, we finally obtain $\langle \rho, n' \rangle \approx_T \langle \rho, n'' \rangle$.

Finally, theorem 3.17 implies the familiar notion of store determinacy for terminating programs. □

Corollary 3.18. If $\langle \sigma, m \rangle \rightarrow_T^* \langle \sigma', \epsilon \rangle$ and $\langle \sigma, m \rangle \rightarrow_T^* \langle \sigma'', \epsilon \rangle$, then $\sigma' = \sigma''$.

Proof. By theorem 3.17. □

4. Sound and Unsound Optimizations

We have achieved our project of proving the essence of trace compilation correct, yet at the same time that result is largely interesting due to its modularity with respect to the O function. In this section we show the example O from section 2.4 to be sound and explore which kinds of O functions are sound and which are not sound.

4.1 Soundness of Variable Folding

Let F, FV , and O be the ones presented in figure 4. For brevity we assume that $FV(s)$ is defined in the usual way and correctly generates the set of free variables for s . That is, for some store σ , $FV(s)$ the set of variables which s never writes to in σ during reduction.

Lemma 4.1. O is sound.

Proof. Assuming we have for some w, w', k such that $w' k \approx_B w k$ for all stores, we want to show that $O(w', \sigma) k \approx_B w k$ for σ .

Our technique will be showing that $O(w', \sigma) k \approx_B w' k$, and then obtaining the desired result via transitivity of \approx_B .

We proceed by case analysis on the O function.

Case: $s = \text{while } b \text{ do } s_1$.

We want to show that

$$(\text{while } b \text{ do } F(s_1, \sigma, FV(s_1))) k \approx_B (\text{while } b \text{ do } s_1) k$$

It suffices to exhibit a bisimulation relation \mathcal{R} such that

$$\mathcal{R}(\sigma, (\text{while } b \text{ do } F(s_1, \sigma, FV(s_1))) k, (\text{while } b \text{ do } s_1) k)$$

Let $s_1' = F(s_1, \sigma, FV(s_1))$. We claim the following relation is a bisimulation relation for any m, ρ, σ' . Note that σ and the loops are

fixed from the assumption.

$$\begin{aligned} \mathcal{R} = & \{(\rho, m, m)\} \\ & \cup \{(\sigma, (\mathbf{while} \ b \ \mathbf{do} \ s'_1) \ k, \\ & \quad (\mathbf{while} \ b \ \mathbf{do} \ s_1) \ k)\} \\ & \cup \{(\sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s'_1 \ \mathbf{while} \ b \ \mathbf{do} \ s'_1) \ k, \\ & \quad (\mathbf{if} \ b \ \mathbf{then} \ (s_1 \ \mathbf{while} \ b \ \mathbf{do} \ s_1)) \ k)\} \\ & \cup \{(\sigma', (F(n, \sigma, FV(s_1)) \ \mathbf{while} \ b \ \mathbf{do} \ s'_1) \ k, \\ & \quad (n \ \mathbf{while} \ b \ \mathbf{do} \ s_1) \ k) \\ & \quad | \sigma'(x) = \sigma(x) \text{ for all free variables in } s_1\} \end{aligned}$$

We proceed by case analysis on the left-side reduction step.

Subcase: The left side and right side are the same.

By lemma 3.3 B is deterministic, both sides step using the same rule, producing the same descendants. But then they are in \mathcal{R} by construction.

$$\begin{array}{ccc} \langle \rho, m \rangle & \mathcal{R} & \langle \rho, m \rangle \\ \downarrow \alpha & & \downarrow \alpha \\ \langle \rho', m' \rangle & \mathcal{R} & \langle \rho', m' \rangle \end{array}$$

Subcase: The left side is $(\mathbf{while} \ b \ \mathbf{do} \ s'_1) \ k$ and the right side is $(\mathbf{while} \ b \ \mathbf{do} \ s_1) \ k$. Both sides reduce by way of While.

Their descendants are in \mathcal{R} by construction.

$$\begin{array}{ccc} \langle \sigma, (\mathbf{while} \ b \ \mathbf{do} \ s'_1) \ k \rangle & \mathcal{R} & \langle \sigma, (\mathbf{while} \ b \ \mathbf{do} \ s_1) \ k \rangle \\ \downarrow \text{While } \tau & & \downarrow \text{While } \tau \\ \langle \sigma', (\mathbf{if} \ b \ \mathbf{then} \ (s'_1 \ \cdots) \ k) \rangle & \mathcal{R} & \langle \sigma', (\mathbf{if} \ b \ \mathbf{then} \ (s_1 \ \cdots) \ k) \rangle \end{array}$$

Subcase: The left side is

$$\langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s'_1 \ \mathbf{while} \ b \ \mathbf{do} \ s'_1)) \ k \rangle$$

and the right side is

$$\langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s_1 \ \mathbf{while} \ b \ \mathbf{do} \ s_1)) \ k \rangle$$

Suppose the left side reduce by IfFalse, then both sides step to k , which is already in \mathcal{R} by way of the first subrelation.

$$\begin{array}{ccc} \langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s'_1 \ \cdots) \ k) \rangle & \mathcal{R} & \langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s_1 \ \cdots) \ k) \rangle \\ \downarrow \text{IfFalse } \tau & & \downarrow \text{IfFalse } \tau \\ \langle \sigma, k \rangle & \mathcal{R} & \langle \sigma, k \rangle \end{array}$$

Subcase: The left side is

$$\langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s'_1 \ \mathbf{while} \ b \ \mathbf{do} \ s'_1)) \ k \rangle$$

and the right side is

$$\langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s_1 \ \mathbf{while} \ b \ \mathbf{do} \ s_1)) \ k \rangle$$

Suppose the left side reduce by IfTrue, we can then fill out the diagram as follows. The descendants are in \mathcal{R} by way of the fourth subrelation, as

$$s'_1 = F(s_1, \sigma, FV(s_1))$$

$$\begin{array}{ccc} \langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s'_1 \ \cdots) \ k) \rangle & \mathcal{R} & \langle \sigma, (\mathbf{if} \ b \ \mathbf{then} \ (s_1 \ \cdots) \ k) \rangle \\ \downarrow \text{IfTrue } \tau & & \downarrow \text{IfTrue } \tau \\ \langle \sigma, (s'_1 \ \cdots) \ k \rangle & \mathcal{R} & \langle \sigma, (s_1 \ \cdots) \ k \rangle \end{array}$$

Subcase: The left side is

$$(F(n, \sigma, FV(s_1)) \ \mathbf{while} \ b \ \mathbf{do} \ s'_1) \ k$$

and the right side is

$$(n \ \mathbf{while} \ b \ \mathbf{do} \ s_1) \ k$$

The left side steps by way of Assign. By inversion, $n = c \ n'$ and

$$F(n, \sigma, FV(s_1)) = F(c, \sigma, FV(s_1)) \ F(n', \sigma, FV(s_1))$$

Further, $c = x := e$. For brevity let $n'' = F(n', \sigma, FV(s_1))$.

By case analysis on e we have two subcases. In the case where $e = n$, we have an identity. In the case where $e = x' + 1 \wedge x' \in v$, $e = \sigma(x') \oplus 1$. We know that $\sigma(x')$ is defined from assumption that the left side steps at all. This means the left side step looks like the following. The call to F is abbreviated due to space.

$$\begin{aligned} \langle \sigma', (c' \ n'' \ \cdots) \ k \rangle & \xrightarrow{\delta}_B \langle \sigma'[x/\sigma(x') \oplus 1], (n'' \ \cdots) \ k \rangle \\ & \text{where } \delta = x/\sigma(x') \oplus 1 \end{aligned}$$

By inversion then we see that the right side, starting with c , also steps by Assign. By the definition of $\hat{\sigma}'$ we have the following reduction for the right side

$$\begin{aligned} \langle \sigma', (c \ n' \ \cdots) \ k \rangle & \xrightarrow{\delta'}_B \langle \sigma'[x/\sigma'(x') \oplus 1], (n' \ \cdots) \ k \rangle \\ & \text{where } \delta' = x/\sigma'(x') \oplus 1 \end{aligned}$$

For these two descendants to be in \mathcal{R} , we need $\sigma'(x') = \sigma(x')$. We know this to hold for all free variables in s_1 , as their freeness guarantees them to be never written to during s_1 's reduction. We assumed that FV correctly generates the set of free variables for a statement. It is easy to see that c is a descendant of s_1 , thus $x \in FV(s_1)$ and $\sigma'(x') = \sigma(x')$ holds.

Since FV is correct, x is not free in s_1 . Therefore,

$$\sigma'[x/\sigma'(x') \oplus 1](y) = \sigma(y)$$

for all y free in s_1 . This finally gives us

$$\mathcal{R}(\sigma'[x/\sigma'(x') \oplus 1], n'', n')$$

which holds by way of the fourth subrelation.

For the diagram below, let $v = \sigma(x') \oplus 1$.

$$\begin{array}{ccc} \langle \sigma', (F(n, \sigma, FV(s_1)) \ \cdots) \ k \rangle & \mathcal{R} & \langle \sigma', (n \ \cdots) \ k \rangle \\ \downarrow \text{Assign } \delta & & \downarrow \text{Assign } \delta \\ \langle \sigma'[x/v], (F(n', \sigma, FV(s_1)) \ \cdots) \ k \rangle & \mathcal{R} & \langle \sigma'[x/v], (n' \ \cdots) \ k \rangle \end{array}$$

All other cases (where s is something other than a **while** loop) are identities. The proofs for the converses are symmetric.

We have shown $O(w', \sigma) \ k \approx_B w' \ k$ for σ . By transitivity, we have the desired result of $O(w', \sigma) \ k \approx_B w \ k$. \square

The most interesting part of the proof is that in every subcase we relied on the right side to be able to mirror the left side's move exactly in a single step. This is a stronger property than required by the bisimulation definition, which says only *visible* moves need to be mirrored.

$$F(b, \sigma, v) = \begin{cases} \text{true} & \text{if } b = x = 0 \wedge x \in v \wedge \sigma(x) = 0 \\ \text{false} & \text{if } b = x = 0 \wedge x \in v \wedge \sigma(x) \neq 0 \\ \text{true} & \text{if } b = x \neq 0 \wedge x \in v \wedge \sigma(x) \neq 0 \\ \text{false} & \text{if } b = x \neq 0 \wedge x \in v \wedge \sigma(x) = 0 \\ \text{undef} & \text{otherwise} \end{cases}$$

$$F(c, \sigma, v) = \begin{cases} \dots \\ \epsilon & \text{if } c = \mathbf{bail } b \text{ to } s_1 \wedge F(b, \sigma, v) = \text{false} \\ \dots \end{cases}$$

Figure 5. Variable Folding extended with Dead Branch Elimination

$$F(c, \sigma, v) = \begin{cases} \dots \\ \epsilon & \text{if } c = x := e \wedge x \text{ has no use sites in the trace} \\ \dots \end{cases}$$

Figure 6. Variable Folding extended with Dead Branch and Dead Store Elimination

4.2 Soundness of Dead Branch Elimination

What kinds of optimizations only mirror visible moves? One can imagine that during tracing we may generate many spurious side-exits. Suppose we extend our variable folding example to also eliminate “dead”, or always-*false* **bails**. The modifications needed for F are shown in figure 5.

For example, considered the following example trace with a dead side-exit. Clearly x is free in the body of the traced loop, and the boolean expression $x \neq 0$ is always going to be *false*.

Example Trace with Dead Bail

```

1  while x = 0 do
2    bail x ≠ 0 to k1
3    z := 1;

```

Plugging the above example into the extended O function will output the following.

Example Trace with Dead Bail Optimized Away

```

1  while x = 0 do
2    z := 1;

```

Such an optimization does not generate code that exactly mirrors the original. This fails to hold if we wholly excise dead conditionals, as the original code would still need to take a step to evaluate the conditional to false before skipping it. To show that this new optimization is still bisimilar, let us extend lemma 4.1 with the proof sketch of a new subcase and its converse.

New subcase and its converse for lemma 4.1.

Subcase: The left side is

$$(F(n, \sigma, FV(s_1)) \mathbf{while } b \mathbf{ do } s'_1) k$$

and the right side is

$$(n \mathbf{while } b \mathbf{ do } s_1) k$$

The left side takes some step. Let $n = c n'$ and

$$F(n, \sigma, FV(s_1)) = F(n', \sigma, FV(s_1))$$

We are concerned with the case when $c = \mathbf{bail } b' \text{ to } k' \wedge F(b', \sigma, FV(s_1)) = \text{false}$, all other cases for c are identities. For brevity let $n'' = F(n', \sigma, FV(s_1))$.

The left side step looks like the following for some n''' . The call to F is abbreviated due to space.

$$\langle \sigma', (n'' \dots) k \rangle \xrightarrow{\alpha} \langle \sigma'', (n''' \dots) k \rangle$$

By inversion we know that $\hat{\sigma}(b') = \text{false}$. Since we assumed that FV correctly generates the set of free variables for s_1 and c is a s_1 -descendant, so $\hat{\sigma}'(b') = \text{false}$. By inversion then we see that the right side, starting with c , steps by **BailFalse**.

$$\langle \sigma', (c n' \dots) k \rangle \xrightarrow{\tau} \langle \sigma', (n' \dots) k \rangle$$

It remains to show that n' can take a step to match the left side step that $F(n', \sigma, FV(s_1))$ took. We again decompose n' into its first command and continuation. We iteratively apply the same reasoning we just underwent until the first command is not **bail** b'' to $k'' \wedge F(b'', \sigma, FV(s_1)) = \text{false}$. For these other cases F acts as an identity for the first command and as congruence for the continuation, so clearly it will take the same α step.

In the diagram below, let $^+$ mean “1 or more times”.

$$\begin{array}{ccc}
\langle \sigma', (F(n', \sigma, FV(s_1)) \dots) k \rangle & \mathcal{R} & \langle \sigma', (n \dots) k \rangle \\
\downarrow \alpha & & \downarrow \tau \text{ BailFalse}^+ \\
& & \langle \sigma', (n' \dots) k \rangle \\
& & \downarrow \alpha \\
\langle \sigma'', (F(n'', \sigma, FV(s_1)) \dots) k \rangle & \mathcal{R} & \langle \sigma'', (n'' \dots) k \rangle
\end{array}$$

The converse is considerably simpler. We have the case where the right side steps by **BailFalse**. By the same reasoning above concerning free variables, we see that the left side would have had its **bail** optimized away into ϵ , thus we can complete the diagram by using *Id*.

$$\begin{array}{ccc}
\langle \sigma', (F(n', \sigma, FV(s_1)) \dots) k \rangle & \mathcal{R} & \langle \sigma', (n \dots) k \rangle \\
\downarrow \text{Id} & & \downarrow \tau \text{ BailFalse} \\
\langle \sigma', (F(n', \sigma, FV(s_1)) \dots) k \rangle & \mathcal{R} & \langle \sigma', (n' \dots) k \rangle
\end{array}$$

All other cases are still identities. \square

4.3 Unsoundness of Dead Store Elimination

Finally, we want to explore what kinds of optimizations are simply unsafe in the tracing framework. Put formally, we want to ask what kind of optimizations do not produce bisimilar code. Continuing with our existing O function, suppose we were to extend it with dead store elimination. That is, suppose variables that we assign to but have no use sites inside the trace body are simply excised. This is shown informally in figure 6.

For example, consider the following example trace with a dead assignment. The variable z is assigned but never used.

Example Trace with Dead Assignment

```
1 while x = 0 do
2   z := 1;
```

Plugging the above example into the extended O function will output the following.

Example Trace with Dead Assignment Optimized Away

```
1 while x = 0 do
2    $\epsilon$ 
```

Intuitively this is unsafe because even though z is dead inside the trace, there very well may be use sites of z after the trace! This intuition is reflected formally. Taking our example above, we need to show that $z := 1$; takes a step that can be mirrored by ϵ . For some σ , by inversion $z := 1$ can step only by Assign : $\langle \sigma, z := 1 \rangle \xrightarrow{\tau} \langle \sigma[z/1], \epsilon \rangle$. ϵ needs to be able to match this move, but $\langle \sigma, \epsilon \rangle \not\xrightarrow{\tau} \langle \sigma[z/1], s \rangle$ for any s . In fact, it does not step at all.

In this fashion this optimization does not output bisimilar code, and is not safe for use inside the tracing framework.

4.4 Soundness of Composition

One property that correct optimizations enjoy in our framework is that the composition of two correct optimizations also yield a correct optimization. We give the following two lemmas to demonstrate this property.

Lemma 4.2. Let $I : (\text{Statement} \times \text{Store}) \rightarrow \text{Statement}$ be the identity function on its first argument. I is sound.

Proof. Trivial by the definition of O -soundness. \square

Lemma 4.3. Let $F, G : (\text{Statement} \times \text{Store}) \rightarrow \text{Statement}$ be two sound optimizations. Let their composition be defined as

$$F \circ G = \lambda(s, \sigma). F(G(s, \sigma), \sigma)$$

$F \circ G$ is sound.

Proof. We want to show that for any w, w', k such that $w' k \approx_B w k$ for all stores, $(F \circ G)(w', \sigma) k \approx_B w k$ for σ .

By soundness of G on w, w', k we know that

$$G(w', \sigma) k \approx_B w k$$

By soundness of F on $w, G(w', \sigma), k$ we then know that

$$F(G(w', \sigma), \sigma) k \approx_B w k$$

But $F(G(w', \sigma), \sigma) k = (F \circ G)(w', \sigma)$, so we are done. \square

5. Related Work

The work carried out in this paper depends on both compiler-correctness and concurrency techniques. Though the corpora of both communities are large, there is a dearth of truly relevant papers that explore purely operational compiler correctness of JIT compilers from as a high-level as ours. Nevertheless, we have taken inspiration as well as fruitful comparisons with several works.

Relevant is Wand’s work on parallel compiler correctness [20]. We believe Wand’s enterprise to be, though also employing bisimulations to prove compiler-correctness, of a different flavor than our own. His approach is the combination of (syntax-directed) denotational semantics and essentially β -convertibility. His picture is also closer to the traditional picture of compiler correctness [5]—that is, the compilation process preserves denotation up to bisimulation—than ours, as his compiler is an ahead-of-time compiler. The elegance of Wand’s work is that he recognized that β -convertibility induces bisimilarity; in the conclusion he admits that almost all the required reasoning is done in the λ -calculus and as such, he can re-use work already done in sequential compiler correctness.

Unlike Wand’s work, our vision of correctness is purely syntax-directed: the translation itself (if the JIT tracing can be seen as such) becomes a non-instantaneous process since we have to spell it out in the operational semantics. This is what makes the enterprise non-trivial. Our notion of convertibility informally becomes something akin to “store-convertibility”, but this is much less powerful than β -convertibility as it does not directly imply bisimilarity. We also do not have the luxury of bringing to bear the entirety of the λ -calculus machinery, so our technique here, while still using bisimulations, is at once more basic and less elegant.

There is also a breadth of literature exploring using bisimulation to show program equivalence by way of contextual equivalence in Pierce et al., Lassen et al., and Wand et al. [10, 17–19]. Though the topics of their specific investigations differ, they all concern themselves with using bisimulation as a more tractable proof technique to prove contextual equivalence without having to universally quantify over all contexts. Their setting is higher-ordered, modeled within the λ -calculus. Pierce et al. [19], for instance, aims to prove bisimilarity sound and complete with respect to contextual equivalence for a modified λ -calculus with recursive types. Wand [10] aims to improve the proof technique and reasons about a λ -calculus extended with explicit stores. These works are basic investigations into the nature of the proof technique. Ours is an application of the technique to prove equivalence of a dynamically transformed program. We also arrived at bisimilarity by an entirely different motivation, that of proving determinism of a JIT compiler that performs the dynamic transform. We are not met with the difficulty of universally quantifying contexts; in fact, we fix our correctness to hold for one context only.

Myreen’s method of creating formally correct JIT compilers for x86 [15] is at the much lower level of abstraction: machine language. They use Hoare logic, and so still retain a flavor of the denotational. We are much farther from the “bare metal” than they are.

6. Conclusions and Future Work

We have demonstrated a paradigm for high-level, purely operational correctness of the tracing JIT compilation technique via bisimulation and confluence. Unlike traditional ahead-of-time compiler correctness where the translation process from the source language to the target language is an opaque function, trace compiler-correctness requires the translation—the tracing—to be spelled out explicitly. We overcome this difficulty by using bisimulations, though we strive to maintain continuity with existing purely operational correctness approaches by returning to confluence.

We hope that the theoretical framework we have provided will prove useful in reasoning about trace compilers at a high level. We hope that we have opened up a wealth of possible future research in the foundational differences between traditional and trace optimizations. Though a different problem, we also feel applying the trace compilation technique to an applicative setting, namely the λ -calculus, will be a worthy venture. It is also interesting to further explore O and the question of just what exactly is observable in computation. We also hope to look at deriving tools from the techniques described here in the future.

Acknowledgements. We thank Michael Bebenita and the Mozilla JavaScript team for enlightening discussions about the implementation of trace compilers. We also thank Jonathan Lee, Oren Freiberger, Kannan Goudan, Dave Herman, Dimitris Vardoulakis, and anonymous reviewers for draft reading and helpful discussions.

References

- [1] Vasanth Bala, Evelyn Duesterwald, and Sanjeev Banerjia. Dynamo: A transparent dynamic optimization system. In *PLDI '00*, pages 1–12. ACM, 2000.
- [2] Michael Bebenita, Florian Brandner, Manuel Fahndrich, Francesco Logozzo, Wolfram Schulte, Nikolai Tillmann, and Herman Venter. SPUR: A trace-based JIT compiler for CIL. In *OOPSLA '10*, 2010.
- [3] Michael Bebenita, Mason Chang, Gregor Wagner, Christian Wimmer, Andreas Gal, and Michael Franz. Trace-based compilation in execution environments without interpreters. In *PPPJ '10*, 2010.
- [4] Mason Chang, Edwin W. Smith, Rick Reitmaier, Michael Bebenita, Andreas Gal, Christian Wimmer, Brendan Eich, and Michael Franz. Tracing for web 3.0: trace compilation for the next generation web applications. In *VEE*, pages 71–80, 2009.
- [5] Joëlle Despeyroux. Proof of translation in natural semantics. In *LICS*, pages 193–205, 1986.
- [6] Cormac Flanagan and Matthias Felleisen. The semantics of future and its use in program optimization. In *POPL '95*, pages 209–220. ACM, 1995.
- [7] Andreas Gal. *Efficient bytecode verification and compilation in a virtual machine*. PhD thesis, 2006. Adviser: Michael Franz.
- [8] Andreas Gal, Brendan Eich, Mike Shaver, David Anderson, David Mandelin, Mohammad R. Haghighat, Blake Kaplan, Graydon Hoare, Boris Zbarsky, Jason Orendorff, Jesse Ruderman, Edwin W. Smith, Rick Reitmaier, Michael Bebenita, Mason Chang, and Michael Franz. Trace-based just-in-time type specialization for dynamic languages. In *PLDI '09*, pages 465–478. ACM, 2009.
- [9] A. J. Kfoury, Michael A. Arbib, and Robert N. Moll. *A Programming Approach to Computability*. Springer-Verlag, 1982.
- [10] Vasileios Koutavas and Mitchell Wand. Small bisimulations for reasoning about higher-order imperative programs. In *POPL '06*, pages 141–152. ACM, 2006.
- [11] David Lacey, Neil D. Jones, Eric Van Wyk, and Carl Christian Fredriksen. Proving correctness of compiler optimizations by temporal logic. In *POPL '02*, pages 283–294. ACM, 2002.
- [12] Sorin Lerner, Todd Millstein, and Craig Chambers. Automatically proving the correctness of compiler optimizations. In *PLDI '03*, pages 220–231. ACM, 2003.
- [13] Mozilla Metrics. Firefox usage: <https://metrics.mozilla.com/>.
- [14] Robin Milner. *Communication and Concurrency*. Prentice Hall, 1995.
- [15] Magnus O. Myreen. Verified just-in-time compiler on x86. In *POPL '10*, pages 107–118. ACM, 2010.
- [16] Frank Pfenning. A proof of the Church-Rosser theorem and its representation in a logical framework. *Journal of Automated Reasoning*, 1993.
- [17] Kristian Støvring and Soren B. Lassen. A complete, co-inductive syntactic theory of sequential control and state. In *POPL '07*, pages 161–172. ACM, 2007.
- [18] Eijiro Sumii and Benjamin C. Pierce. A bisimulation for dynamic sealing. In *POPL '04*, pages 161–172. ACM, 2004.
- [19] Eijiro Sumii and Benjamin C. Pierce. A bisimulation for type abstraction and recursion. In *POPL '05*, pages 63–74. ACM, 2005.
- [20] Mitchell Wand. Compiler correctness for parallel languages. In *FPCA*, pages 120–134, 1995.
- [21] Mitchell Wand and William D. Clinger. Set constraints for destructive array update optimization. *Journal of Functional Programming*, 11(3):319–346, 2001.
- [22] Mitchell Wand and Igor Siveroni. Constraint systems for useless variable elimination. In *POPL '99*, pages 291–302. ACM, 1999.