

Proving Acceptability Properties of Relaxed Nondeterministic Approximate Programs

Michael Carbin Deokhwan Kim Sasa Misailovic Martin C. Rinard

MIT CSAIL

{mcarbin, dkim, misailo, rinard}@csail.mit.edu

Abstract

Approximate program transformations such as task skipping [29, 30], loop perforation [21, 22, 34], multiple selectable implementations [3, 4, 16], approximate function memoization [11], and approximate data types [33] produce programs that can execute at a variety of points in an underlying performance versus accuracy tradeoff space. Namely, these transformed programs trade accuracy of their results for increased performance by dynamically and nondeterministically modifying variables that control their execution.

We call such transformed programs *relaxed* programs—they have been extended with additional nondeterminism to *relax* their semantics and enable greater flexibility in their execution.

In this paper, we present programming language constructs for developing and specifying relaxed programs. We also present proof rules for reasoning about *acceptability properties* [28], which are the key correctness properties of relaxed programs that enable them to have a range of acceptable executions. Specifically, our proof rules enable programmers to directly specify and verify *relational properties* that characterize the desired correctness relationships, such as accuracy, between the values of variables in a program’s original semantics (before the transformation) and its relaxed semantics. Our proof rules also support verifying *safety properties* which characterize desirable properties involving values in the current execution and must hold for all acceptable executions.

The rules are designed to support a reasoning approach in which the majority of the reasoning effort uses the original semantics. This effort is then transferred to establish the desired properties of the program under the relaxed semantics by using relational reasoning to bridge the gap between the two semantics.

We have formalized the dynamic semantics of our target programming language and the proof rules in Coq, and verified that the proof rules are sound with respect to the dynamic semantics. Our Coq implementation enables developers to obtain fully machine checked verifications of their relaxed programs.

Categories and Subject Descriptors D.3.2 [Programming Languages]: Language Classifications—Nondeterministic languages; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs

General Terms Languages, Performance, Theory, Verification

Keywords Coq, Relational Hoare Logic, Relaxed Programming

1. Introduction

In recent years researchers have developed a range of mechanisms for dynamically varying application behavior. Typical goals include maximizing performance subject to an accuracy constraint, maximizing accuracy subject to a performance constraint, or dynamically adjusting program behavior to adapt to changes in the characteristics of the underlying hardware platform (such as varying load or clock rate) [16, 17]. Specific mechanisms include multiple selectable implementations of a given component or components [3, 4, 37]; sampling inputs to reductions [37], skipping tasks in parallel programs [29, 30]; loop perforation (skipping iterations of time-consuming loops) [21, 22, 34]; approximate function memoization (returning a previously computed value when the arguments of the new function call are close to the arguments of the previous function call) [11]; dynamic knobs (configuration parameters that can be changed as the program executes) [16]; and approximate data types (data types that return approximate results) [33].

All of these mechanisms can produce a *relaxed* program—a program that may adjust its execution by changing one or more variables in the program, subject to a specified constraint or relaxation predicate. For example, a program transformed with loop perforation may dynamically choose to skip loop iterations each time it enters a loop. A relaxed program is therefore a nondeterministic program, with each execution a variant of the original execution. The different executions typically share a common global structure, with local differences at only those parts of the computation that are affected by the change.

1.1 Reasoning About Relaxed Programs

In this paper, we present a language for writing relaxed programs and proof rules for reasoning about the acceptability properties of relaxed programs. The language and proof rules are designed to support a staged approach in which the developer first develops a standard program and uses standard approaches to reason about the program to verify that it satisfies desired correctness properties. We refer to the dynamic semantics of the program at this stage as the *original semantics* of the program.

Either the developer or an automated system (such as a compiler that implements loop perforation) then *relaxes* the program to enable the desired additional nondeterministic executions. We refer to the dynamic semantics of the program at this stage as the *relaxed semantics* of the program.

Finally, the developer verifies that the relaxation maintains the desired correctness properties by augmenting her original reasoning with new reasoning that relates the relaxed semantics to the original semantics. Specifically, it is now possible for the developer to specify and verify additional *relational properties* that relate the relaxed semantics to the original semantics. This approach is designed to reduce the overall reasoning effort by exploiting the structure that the original and relaxed programs share. With this ap-

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.

PLDI’12, June 11–16, 2012, Beijing, China.

Copyright © 2012 ACM 978-1-4503-1205-9/12/06...\$10.00

proach the majority of the reasoning effort works with the original program semantics and is then transferred via relational reasoning to verify the nondeterministic relaxed program.

1.2 Relaxed Programming Constructs

The basic programming constructs of relaxed programming include nondeterministic assignment to variables (via the `relax` statement); relational properties that relate the relaxed semantics to the original semantics (via the `relate` statement); assertions (via the `assert` statement); and assumptions (via the `assume` statement).

The Relax Statement. The `relax (X) st (P)` statement specifies a nondeterministic assignment to the set of variables X . Specifically, the `relax` statement can assign the variables in X to any set of values that satisfies the relaxation predicate P . When added to the program, the `relax` statement affects only the relaxed semantics of the program; in the original semantics it has no effect.

The Relate Statement. The `relate P` statement asserts that the property P must hold at the program point where the statement appears. The property P in `relate` statements may reference values from both the original and relaxed semantics. So, for example, the statement might require that the value of a variable x in the relaxed semantics must be greater than or equal to the value of that variable in the original semantics. Such properties are relational properties because they relate states from the two semantics.

The Assert Statement. The `assert P` statement states that P must hold at the point where the statement appears. In contrast to the `relate` statement, P only references values from a single semantics (original or relaxed). In the original semantics, our proof rules generate an obligation to prove that an `assert` statement holds for all executions. To ensure that the `assert` statement remains valid in the relaxed semantics, our proof rule in the relaxed semantics generates an obligation to prove that given the validity of the assertion in the original semantics, then the current relation between the original and relaxed semantics establishes that the assertion is valid in the relaxed semantics. For example, it may be possible to prove that all the variables referenced in an assertion have the same values in the original and relaxed semantics—namely, relaxation does not interfere with the assertion. In this way, relational reasoning serves as a bridge to transfer properties of the original program over to the relaxed program.

The Assume Statement. The `assume P` statement states that the property P holds at the point where the statement appears. In the original semantics the `assume` statement does not generate any proof obligations—the proof system simply accepts that P holds. However, because relaxation may interfere with the reasoning behind the assumption, the proof rules for the relaxed semantics must generate an obligation to prove the assumption valid. The reasoning works much in the same way as the `assert` statement except that our proof rules prove that `assert` statements are valid in the original semantics whereas `assume` statements are assumed valid.

1.3 Key Considerations for Acceptable Relaxed Programs

Our approach formalizes several considerations that are critical for the development and deployment of acceptable relaxed programs:

- **Safety:** Most programs have critical functional correctness properties, or *safety properties*, that must hold for all executions of the program and, however, may be compromised by relaxation. The `assert` statement ensures the developers can specify and verify properties that must hold for both original and relaxed executions of the program.
- **Accuracy:** In addition to critical functional correctness properties, many programs have soft functional correctness properties or *accuracy properties* that characterize what makes the

program amenable to relaxation. For example, loops that are amenable to loop perforation can often return a range of acceptable results—as opposed to a single, absolutely correct result. The `relate` statement enables developers to specify accuracy properties, which are inherently relational, by for example bounding the difference between the outputs of the original and relaxed program.

- **Debugging:** Like the vast majority of reasoning systems, the `assume` statement provides developers with the ability to use unverified assumptions to ease or speed the development and verification of their programs. However, if a developer’s assumptions are not true, then the program may exhibit bugs or unintended behaviors.

Relaxation makes reasoning about unintended behaviors and, in particular, debugging unintended behaviors more difficult because there are now additional reasons why a program could behave unexpectedly: it may be the case that relaxation introduced the unintended behavior, or that relaxation interfered with the developer’s pre-existing assumptions. But, it is also possible that the developer’s assumptions were not true and, hence, the original program contained the unintended behavior. Our proof rules ensure that assumptions about the original program, if valid, are still valid for the relaxed program.

Note that our design enables these statements to work together to prove important properties. For example, the developer may use a `relate` statement to establish a relationship between values in variables from original and relaxed executions, then use this relationship to prove that a given property in an `assert` statement holds in the relaxed execution.

1.4 Proof Rules and Formal Properties

We structure our program logic for relaxed programs as set of Hoare logic proof rules:

- **Axiomatic Original Semantics:** The Hoare-style semantics models the original execution of the program wherein `relax` statements have no effect.
- **Axiomatic Relaxed Semantics:** The *relational*, Hoare-style semantics relates executions in the relaxed semantics to executions in the original semantics. The predicates of the judgment are given in a relational logic that enables us to express properties over the value of variables in both the original and relaxed execution of the program. A proof with the axiomatic relaxed semantics relates the two semantics of the program in lockstep and, therefore, supports the transfer of reasoning about the original semantics by enabling, for example, noninterference proofs that show that the values of the variables in the relaxed execution are the same as those in the original execution.

One key aspect of the axiomatic relaxed semantics is that it must also give an appropriate semantics for relaxed programs in which original and relaxed executions may branch in different directions at a control flow construct because then the two executions are no longer in lockstep. In particular, relational reasoning cannot be used for control flow constructs where this may happen and, more importantly, `relate` statements under these constructs do not have a natural semantics. To address this, our relaxed semantics incorporates a non-relational axiomatic *intermediate* semantics that captures the desired behavior of the relaxed execution as it executes without a corresponding original execution—i.e., it does not violate pre-existing assumptions—and restricts the location of `relate` statements.

Our proof rules are sound and establish the following semantic properties of relaxed programs:

- **Original Progress Modulo Assumptions:** If the program verifies under the axiomatic original semantics, then no execution of the program in the dynamic original semantics violates an assertion. By design, a program may still terminate in error if a specified assume statement is not valid.
- **Relative Relaxed Progress:** If the program verifies under the axiomatic relaxed semantics and no executions in the dynamic original semantics violate an assertion or an assumption, then no execution in the dynamic relaxed semantics violates an assertion or an assumption.
- **Relaxed Progress:** If the program verifies under both the original and relaxed axiomatic semantics and no execution in the dynamic original semantics violates an assumption, then no execution in the dynamic relaxed semantics violates an assertion or an assumption.
- **Relaxed Progress Modulo Original Assumptions:** If the program verifies under both the original and relaxed axiomatic semantics, then if an execution in the dynamic relaxed semantics violates an assertion or an assumption, then an execution in the dynamic original semantics violates an assumption.
- **Relational Soundness:** If the program verifies under the axiomatic relaxed semantics, then all pairs of executions in the dynamic original and relaxed semantics satisfy the `relate` statements in the program.

1.5 Coq Verification Framework

We have formalized the dynamic original and relaxed semantics with the Coq proof assistant [1]. We have also used Coq to formalize the proof rules for the static semantics and obtain a fully machine checked proof that the rules are sound with respect to the dynamic semantics and provide the stated semantic properties. Our Coq formalization makes it possible to develop fully machine checked verifications of relaxed programs. We have used our framework to develop and verify several small relaxed programs.

Our Coq implementation contains approximately 8000 lines of code and proof scripts, with 1300 lines devoted to the original semantics and its soundness proofs and 1900 additional lines devoted to the relaxed semantics and its soundness proofs. A large portion of the implementation (approximately 3500 lines) is devoted to formalizing the semantics of our relational assertion logic and its soundness with respect to operations such as substitution.

1.6 Contributions

This paper makes the following contributions:

- **Relaxed Programming:** It identifies the concept of relaxed programming as a way to specify nondeterministic variants of a program. The variants often occupy a range of points in an underlying performance versus accuracy trade-off space. Current techniques that can produce relaxed programs include multiple selectable implementations [3, 4], skipping tasks [29, 30], loop perforation [21, 22, 34], approximate function memoization [11], approximate data types [33], and dynamic knobs [16].
- **Reasoning Approach and Proof Rules:** It identifies the basic reasoning approach and proof rules for first verifying important safety properties of the original program, stating relational properties that may reference the states of both the original and relaxed programs, then augmenting the verification of the original program to establish the correctness properties in the relaxed program. With this approach, the majority of the reasoning effort works with the original program and is transferred to the relaxed program through relational reasoning.

```

iop ::= + | - | * | / | ...
cmp ::= < | > | = | ...
lop ::= ^ | v | ...
X ::= x | x, X
E ::= n | x | E iop E
E* ::= n | x(o) | x(r) | E* iop E*
B ::= true | false | E cmp E | B lop B | ¬B
B* ::= true | false | E* cmp E* | B* lop B* | ¬B*
S ::= skip | x = E | havoc (X) st (B) | relax (X) st (B)
    | if (B) {S1} else {S2} | while (B) {S}
    | assume B | assert B | relate l : B*
    | S ; S

```

Figure 1. Language Syntax

- **Coq Formalization and Soundness Results:** It presents a formalization of the dynamic semantics and proof rules in Coq. We have used this formalization to prove that the proof rules are sound with respect to the dynamic original and relaxed semantics. We note that our Coq formalization contains a reusable implementation of our relational assertion logic that is, in principle, suitable for other uses such as verifying traditional compiler transformations [7, 31, 35].
- **Verified Programs:** It presents several relaxed programs for which we have used the Coq formalization to develop fully machine checked verifications.

Relaxed programs can deliver substantial flexibility, performance, and resource consumption benefits. But to successfully deploy relaxed programs, developers need to know that the relaxation preserves important correctness properties and produces an acceptable program. This paper presents a foundational formal reasoning system that leverages the structure and relationships that the original and relaxed executions share to enable the verification of these properties.

2. Language Syntax and Dynamic Semantics

Figure 1 presents a simple imperative language with integer variables, integer arithmetic expressions, boolean expressions, conditional statements, while loops, and sequential composition. For generality, we support nondeterminism in the original semantics via the `havoc (X) st (B)` statement which nondeterministically assigns the variables in `X` to values that satisfy `B`. The `relax (X) st (B)` statement supports nondeterministic relaxation — in the original semantics it has no effect; in the relaxed semantics it nondeterministically assigns the variables in `X` to values that satisfy `B`. The language also supports the standard `assume` and `assert` statements.

A main point of departure from standard languages is the addition of *relational integer expressions* (E^*) and *relational boolean expressions* (B^*). Unlike standard expressions, which involve values from only the current execution, relational expressions can reference values from both the original ($x(o)$) and relaxed ($x(r)$) executions. These relational expressions enable `relate` statements to specify relationships that must hold between the original and relaxed executions. For example, the statement `relate l : x(o) = x(r)` asserts that at the current program point (labeled by `l`), `x` must have the same value in both executions.

$$\begin{aligned}
\text{Vars} &\subseteq \{x, y, z, \dots\} & \llbracket \text{iop} \rrbracket &\in \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{Z} \\
\mathbb{Z} &= \{\dots, -1, 0, 1, \dots\} & \llbracket \text{cmp} \rrbracket &\in \mathbb{Z} \times \mathbb{Z} \rightarrow \mathbb{B} \\
\mathbb{B} &= \{\text{true}, \text{false}\} & \llbracket \text{lop} \rrbracket &\in \mathbb{B} \times \mathbb{B} \rightarrow \mathbb{B} \\
\sigma \in \Sigma &= \text{Vars} \xrightarrow{\text{fin}} \mathbb{Z}
\end{aligned}$$

$$\begin{aligned}
\llbracket E \rrbracket &\in \Sigma \rightarrow \mathbb{Z} \\
\llbracket \mathbf{n} \rrbracket(\sigma) &= n \\
\llbracket \mathbf{x} \rrbracket(\sigma) &= \sigma(x) \\
\llbracket E_1 \text{ iop } E_2 \rrbracket(\sigma) &= \llbracket E_1 \rrbracket(\sigma) \text{ iop } \llbracket E_2 \rrbracket(\sigma) \\
\llbracket E^* \rrbracket &\in \Sigma \times \Sigma \rightarrow \mathbb{Z} \\
\llbracket \mathbf{n} \rrbracket(\sigma_1, \sigma_2) &= n \\
\llbracket \mathbf{x}(\mathbf{o}) \rrbracket(\sigma_1, \sigma_2) &= \sigma_1(\mathbf{x}) \\
\llbracket \mathbf{x}(\mathbf{r}) \rrbracket(\sigma_1, \sigma_2) &= \sigma_2(\mathbf{x}) \\
\llbracket E_1^* \text{ iop } E_2^* \rrbracket(\sigma_1, \sigma_2) &= \llbracket E_1^* \rrbracket(\sigma_1, \sigma_2) \text{ iop } \llbracket E_2^* \rrbracket(\sigma_1, \sigma_2) \\
\llbracket B \rrbracket &\in \Sigma \rightarrow \mathbb{B} \\
\llbracket \text{true} \rrbracket(\sigma) &= \text{true} \\
\llbracket \text{false} \rrbracket(\sigma) &= \text{false} \\
\llbracket E_1 \text{ cmp } E_2 \rrbracket(\sigma) &= \llbracket E_1 \rrbracket(\sigma) \text{ cmp } \llbracket E_2 \rrbracket(\sigma) \\
\llbracket B_1 \text{ lop } B_2 \rrbracket(\sigma) &= \llbracket B_1 \rrbracket(\sigma) \text{ lop } \llbracket B_2 \rrbracket(\sigma) \\
\llbracket \neg B \rrbracket(\sigma) &= \begin{cases} \text{true}, & \llbracket B \rrbracket(\sigma) = \text{false} \\ \text{false}, & \llbracket B \rrbracket(\sigma) = \text{true} \end{cases} \\
\llbracket B^* \rrbracket &\in \Sigma \times \Sigma \rightarrow \mathbb{B} \\
\llbracket \text{true} \rrbracket(\sigma_1, \sigma_2) &= \text{true} \\
\llbracket \text{false} \rrbracket(\sigma_1, \sigma_2) &= \text{false} \\
\llbracket E_1^* \text{ cmp } E_2^* \rrbracket(\sigma_1, \sigma_2) &= \llbracket E_1^* \rrbracket(\sigma_1, \sigma_2) \text{ cmp } \llbracket E_2^* \rrbracket(\sigma_1, \sigma_2) \\
\llbracket B_1^* \text{ lop } B_2^* \rrbracket(\sigma_1, \sigma_2) &= \llbracket B_1^* \rrbracket(\sigma_1, \sigma_2) \text{ lop } \llbracket B_2^* \rrbracket(\sigma_1, \sigma_2) \\
\llbracket \neg B^* \rrbracket(\sigma_1, \sigma_2) &= \begin{cases} \text{true}, & \llbracket B^* \rrbracket(\sigma_1, \sigma_2) = \text{false} \\ \text{false}, & \llbracket B^* \rrbracket(\sigma_1, \sigma_2) = \text{true} \end{cases} \\
\text{expr} \frac{\llbracket e \rrbracket(\sigma) = n}{\langle e, \sigma \rangle \downarrow_E n} & \quad \text{bexp} \frac{\llbracket b \rrbracket(\sigma) = v}{\langle b, \sigma \rangle \downarrow_B v}
\end{aligned}$$

Figure 2. Semantics of Expressions

2.1 Semantics of Expressions

Figure 2 presents the semantics of expressions in the language. The denotations of expressions are functions mapping a state or pair of states to either an integer (\mathbb{Z}) or boolean value (\mathbb{B}). A state (σ) is a finite map from program variables, Vars , to integers, \mathbb{Z} , and is an element of the domain Σ , which is the set of all finite maps from variables to integers.

The semantic function $\llbracket E \rrbracket$ defines the semantics for integer expressions, which are composed of the standard integer operations (e.g., $+$, $-$, $*$, $/$, ...) on integer operands. The semantic function $\llbracket B \rrbracket$ defines the semantics of boolean expressions, which are composed of the standard comparison operators on integers (e.g., $<$, $=$, $>$, ...) and the standard boolean operators (e.g., \wedge , \vee , ...).

The semantic function $\llbracket E^* \rrbracket$ defines the semantics for relational integer expressions as a function mapping a pair of states (σ_1, σ_2) to an integer number. Our convention is to have the first component of the state pair be a state from the original semantics and the

second component a state from the relaxed semantics. Therefore, a reference to a variable in the original semantics, $\mathbf{x}(\mathbf{o})$, is equivalent to $\sigma_1(x)$ whereas a reference to a variable, $\mathbf{x}(\mathbf{r})$, in the relaxed semantics is equivalent to $\sigma_2(x)$.

The semantic function $\llbracket B^* \rrbracket$ likewise extends the semantics for boolean expressions with the capability to express boolean properties over relational integer expressions.

2.2 Dynamic Original Semantics

Figure 3 presents the dynamic original semantics of the program in a big-step operational style. The evaluation relation $\langle s, \sigma \rangle \Downarrow_o \phi$ means that evaluation of the statement s from a state σ yields the output configuration ϕ . An output configuration is an element in the domain $\Phi : \text{ba} \cup \text{wr} \cup (\Sigma \times \Psi)$.

The distinguished element ba (“bad assume”) is intended to mean that the program has failed at an `assume` statement in the program. The distinguished element wr (“wrong”) denotes that the program has failed due to another error, such as an unsatisfied `assert` statement in the program.

An element in the domain $\Sigma \times \Psi$ indicates that the program has terminated successfully, yielding a final state σ and an *observation list*, $\psi \in \Psi$, which is the sequence of *observations* emitted by `relate` statements during the execution of the program.

An observation (l, σ) is an element in the domain $L \times \Sigma$. L is the finite domain consisting of all the labels specified in `relate` statements in the program — the execution of each `relate` statement emits an observation consisting of its label along with the current state of the program.

The structure of an observation list is given by the standard constructors for lists: $\Psi = \emptyset \mid (l, \sigma) :: \Psi$. We also use the notation $\psi_1.\psi_2$ to denote the result of appending two lists.

Evaluation Rules

The rules for `skip`, `assignment`, `if`, `sequential composition`, and `while` statements follow the standard semantics for these constructs. The semantics for non-standard constructs are as follows:

Havoc. The `havoc` (X) `st` (e) statement non-deterministically assigns values to the set of variables in X such that their values satisfy the statement’s constraint, e . All variables not specified in X retain their previous values. If there does not exist an assignment of values to X that satisfy e , then the statement evaluates to wr .

Assert. The `assert` e statement checks that the state satisfies its constraint e . If the boolean expression evaluates to true , then evaluation continues; otherwise, the statement evaluates to wr .

Assume. The `assume` e statement checks that the state satisfies its constraint e . If the boolean expression evaluates to true , then evaluation continues; otherwise, the statement evaluates to ba .

Relax. The `relax` (X) `st` (e) statement does not modify the state of the program in the original semantics. However, our design for relaxation is that a `relax` statement strictly relaxes or generalizes the semantics of a program at a given point. In particular, an original execution of the program is a trivial member of the relaxed semantics of the program. Therefore, the dynamic semantics of the `relax` statement assert that the relaxation predicate holds. We implement this by reusing the evaluation rule for `assert`.

Relate. The `relate` $l : e^*$ statement is a relational assertion over original and relaxed executions of the program. As already discussed, the dynamic semantics of the statement is to emit an observation consisting of the statement’s label l along with the current state of the program. This semantics enables us to define and verify a relation on the observation lists emitted by the original and relaxed variants of the program (Section 4).

$$\boxed{\langle s, \sigma \rangle \Downarrow_o \phi}$$

$$\begin{array}{c}
\text{skip} \frac{}{\langle \text{skip}, \sigma \rangle \Downarrow_o \langle \sigma, \emptyset \rangle} \quad \text{assign} \frac{\langle e, \sigma \rangle \Downarrow_E n}{\langle x = e, \sigma \rangle \Downarrow_o \langle \sigma[x \mapsto n], \emptyset \rangle} \quad \text{havoc-t} \frac{\langle e, \sigma' \rangle \Downarrow_B \text{true} \quad \forall_{x \notin X} \cdot \sigma(x) = \sigma'(x)}{\langle \text{havoc}(X) \text{ st } (e), \sigma \rangle \Downarrow_o \langle \sigma', \emptyset \rangle} \\
\text{havoc-f} \frac{\neg \exists \sigma' \cdot (\langle e, \sigma' \rangle \Downarrow_B \text{true} \wedge \forall_{x \notin X} \cdot \sigma(x) = \sigma'(x))}{\langle \text{havoc}(X) \text{ st } (e), \sigma \rangle \Downarrow_o wr} \quad \text{assert-t} \frac{\langle e, \sigma \rangle \Downarrow_B \text{true}}{\langle \text{assert } e, \sigma \rangle \Downarrow_o \langle \sigma, \emptyset \rangle} \quad \text{assert-f} \frac{\langle e, \sigma \rangle \Downarrow_B \text{false}}{\langle \text{assert } e, \sigma \rangle \Downarrow_o wr} \\
\text{assume-t} \frac{\langle e, \sigma \rangle \Downarrow_B \text{true}}{\langle \text{assume } e, \sigma \rangle \Downarrow_o \langle \sigma, \emptyset \rangle} \quad \text{assume-f} \frac{\langle e, \sigma \rangle \Downarrow_B \text{false}}{\langle \text{assume } e, \sigma \rangle \Downarrow_o ba} \quad \text{relax} \frac{\langle \text{assert } e, \sigma \rangle \Downarrow_o \phi}{\langle \text{relax}(X) \text{ st } (e), \sigma \rangle \Downarrow_o \phi} \\
\text{relate} \frac{}{\langle \text{relate } l: e^*, \sigma \rangle \Downarrow_o \langle \sigma, (l, \sigma) \rangle} \quad \text{if-t} \frac{\langle b, \sigma \rangle \Downarrow_B \text{true} \quad \langle s_1, \sigma \rangle \Downarrow_o \phi}{\langle \text{if}(b) \{s_1\} \text{ else } \{s_2\}, \sigma \rangle \Downarrow_o \phi} \quad \text{if-f} \frac{\langle b, \sigma \rangle \Downarrow_B \text{false} \quad \langle s_2, \sigma \rangle \Downarrow_o \phi}{\langle \text{if}(b) \{s_1\} \text{ else } \{s_2\}, \sigma \rangle \Downarrow_o \phi} \\
\text{seq-l} \frac{\langle s_1, \sigma \rangle \Downarrow_o \langle \sigma', \psi_1 \rangle \quad \langle s_2, \sigma' \rangle \Downarrow_o \langle \sigma'', \psi_2 \rangle}{\langle s_1 ; s_2, \sigma \rangle \Downarrow_o \langle \sigma'', \psi_2.\psi_1 \rangle} \quad \text{while-f} \frac{\langle b, \sigma \rangle \Downarrow_B \text{false}}{\langle \text{while}(b) \{s\}, \sigma \rangle \Downarrow_o \langle \sigma, \emptyset \rangle} \\
\text{while-tl} \frac{\langle b, \sigma \rangle \Downarrow_B \text{true} \quad \langle s, \sigma \rangle \Downarrow_o \langle \sigma', \psi_1 \rangle \quad \langle \text{while}(b) \{s\}, \sigma' \rangle \Downarrow_o \langle \sigma'', \psi_2 \rangle}{\langle \text{while}(b) \{s\}, \sigma \rangle \Downarrow_o \langle \sigma'', \psi_2.\psi_1 \rangle}
\end{array}$$

Figure 3. Dynamic Original Semantics

$$\boxed{\langle s, \sigma \rangle \Downarrow_r \phi}$$

$$\begin{array}{c}
\text{relax} \frac{\langle e, \sigma \rangle \Downarrow_B \text{false}}{\langle \text{relax}(X) \text{ st } (e), \sigma \rangle \Downarrow_r wr} \\
\text{relax} \frac{\langle \text{havoc}(X) \text{ st } (e), \sigma \rangle \Downarrow_r \phi}{\langle \text{relax}(X) \text{ st } (e), \sigma \rangle \Downarrow_r \phi}
\end{array}$$

Figure 4. Dynamic Relaxed Semantics

We have elided a presentation of the standard rules for propagating errors (*ba* and *wr*) through the program. However, the reader can refer to our previous technical report for further details [9].

2.3 Dynamic Relaxed Semantics

Figure 4 presents an abbreviated definition of the dynamic relaxed semantics of a program. The evaluation relation $\langle s, \sigma \rangle \Downarrow_r \phi$ means that evaluation of the statement s from a state σ yields the output configuration ϕ .

The dynamic relaxed semantics builds upon the original semantics and differs from the original semantics only in that `relax` statements modify the state of the program. We have, therefore, elided the presentation of all the rules that are either reused (`skip`, `assignment`, `havoc`, `assert`, and `assume`) or adapted to refer to the relaxed dynamic semantics in their premises (i.e., sequential composition, `if`, and `while`).

Relax. The `relax (X) st (e)` statement non-deterministically modifies the state of the program in the relaxed semantics so that it satisfies the statement’s constraint e . The rule implements the modification by reusing the rule for `havoc` statements.

3. Axiomatic Semantics

In this section we present axiomatic definitions for programs in our model. We formalize this problem by presenting axiomatic definitions for the original semantics and the relaxed semantics.

- **Axiomatic Original Semantics.** The proof rules model the dynamic original semantics of the program. If the program verifies with these rules, then no execution of the program in the dynamic original semantics violates an assertion (i.e., evaluates to *wr*). However, the program may dynamically violate an assumption (i.e., evaluate to *ba*).

$$\begin{array}{l}
P ::= \text{true} \mid \text{false} \mid E \text{ cmp } E \mid P \text{ lop } P \mid \neg P \mid \exists_{\mathbf{x}} \cdot P \\
P^* ::= \text{true} \mid \text{false} \mid E^* \text{ cmp } E^* \mid P^* \text{ lop } P^* \mid \neg P^* \\
\quad \mid \exists_{\mathbf{x}(\circ)} \cdot P^* \mid \exists_{\mathbf{x}(\mathbf{r})} \cdot P^*
\end{array}$$

Figure 5. Relational Assertion Logic Syntax

- **Axiomatic Relaxed Semantics.** The proof rules model pairs of executions of the program in the dynamic original and dynamic relaxed semantics. If the program verifies with these rules, then if all executions in the original semantics execute without error (i.e., do not evaluate to *wr* or *ba*), then all executions in the relaxed semantics execute without error. A proof with these rules also guarantees that pairs of original and relaxed executions satisfy all of the `relate` statements in the program.

3.1 Relational Assertion Logic

Figure 5 presents the concrete syntax of our *relational* assertion logic. This logic extends a non-relational assertion logic with relational formulas, which then allow us to reason about the validity of relational boolean expressions in `relate` statements. Its presentation follows Benton’s style in his Relational Hoare Logic [7].

3.1.1 Syntax

The syntactic category P gives the syntax for formulas in first-order logic with integer expressions and existential quantification. The syntactic category P^* gives corresponding syntax for writing relational formulas. P^* extends P by allowing formula to refer to relational integer expressions.

3.1.2 Semantics

Figure 6 presents the semantics of formulas in the logic. The denotation of a non-relational formula $\llbracket P \rrbracket$ is the set of states that satisfy the formula. $\llbracket P \rrbracket$ reuses the semantic definitions for integer expressions from Figure 2 to construct a definition for each formula. The denotation of a relational formula $\llbracket P^* \rrbracket$ closely follows that of non-relational formulas: it is the set of pairs of states that satisfy the relation. In our use of the logic, references to the original semantics (e.g., $\mathbf{x}(\circ)$) refer to the first component of the pair and references to the relaxed semantics (e.g., $\mathbf{x}(\mathbf{r})$) refer to the second component.

$$\begin{aligned}
\llbracket P \rrbracket &\in \mathcal{P}(\Sigma) \\
\llbracket \text{true} \rrbracket &= \Sigma \\
\llbracket \text{false} \rrbracket &= \emptyset \\
\llbracket E_1 \text{ cmp } E_2 \rrbracket &= \{\sigma \mid \llbracket E_1 \rrbracket(\sigma) \text{ cmp } \llbracket E_2 \rrbracket(\sigma)\} \\
\llbracket P_1 \text{ lop } P_2 \rrbracket &= \{\sigma \mid \sigma \in \llbracket P_1 \rrbracket \text{ lop } \sigma \in \llbracket P_2 \rrbracket\} \\
\llbracket \neg P \rrbracket &= \llbracket \text{true} \rrbracket \setminus \llbracket P \rrbracket \\
\llbracket \exists x \cdot P \rrbracket &= \{\sigma \mid n \in \mathbb{Z}, \sigma \in \llbracket P[n/x] \rrbracket\} \\
\llbracket P^* \rrbracket &\in \mathcal{P}(\Sigma \times \Sigma) \\
\llbracket \text{true} \rrbracket &= \Sigma \times \Sigma \\
\llbracket \text{false} \rrbracket &= \emptyset \\
\llbracket E_1^* \text{ cmp } E_2^* \rrbracket &= \{(\sigma_1, \sigma_2) \mid \llbracket E_1^* \rrbracket(\sigma_1, \sigma_2) \text{ cmp } \llbracket E_2^* \rrbracket(\sigma_1, \sigma_2)\} \\
\llbracket P_1^* \text{ lop } P_2^* \rrbracket &= \{(\sigma_1, \sigma_2) \mid (\sigma_1, \sigma_2) \in \llbracket P_1^* \rrbracket \text{ lop } (\sigma_1, \sigma_2) \in \llbracket P_2^* \rrbracket\} \\
\llbracket \neg P^* \rrbracket &= \llbracket \text{true} \rrbracket \setminus \llbracket P^* \rrbracket \\
\llbracket \exists x(\circ) \cdot P^* \rrbracket &= \{(\sigma_1, \sigma_2) \mid n \in \mathbb{Z}, (\sigma_1, \sigma_2) \in \llbracket P^*[n/x(\circ)] \rrbracket\} \\
\llbracket \exists x(\mathbf{r}) \cdot P^* \rrbracket &= \{(\sigma_1, \sigma_2) \mid n \in \mathbb{Z}, (\sigma_1, \sigma_2) \in \llbracket P^*[n/x(\mathbf{r})] \rrbracket\}
\end{aligned}$$

Figure 6. Relational Assertion Logic Semantics

Injections. We also define the injection functions $inj_o(P)$ and $inj_r(P)$, which construct a relational formula in P^* from a non-relational formula in P . Conceptually, the form $inj_o(P)$ constructs a relational formula where P holds for the original semantics by replacing variables (e.g., x) in P with the relational original variables (e.g., $x(\circ)$); $inj_r(P)$ does the same with the relational relaxed variables such that P holds for the relaxed semantics. This means that $inj_o(P)$ (resp. $inj_r(P)$) creates a formula representing all state pairs where the first (resp. second) component satisfies P :

$$\begin{aligned}
\llbracket inj_o(P) \rrbracket &= \{(\sigma_1, \sigma_2) \mid \sigma_1 \in \llbracket P \rrbracket\} \\
\llbracket inj_r(P) \rrbracket &= \{(\sigma_1, \sigma_2) \mid \sigma_2 \in \llbracket P \rrbracket\}
\end{aligned}$$

We also define the notation $\langle P_1 \cdot P_2 \rangle$ for combining a predicate P_1 over the original semantics with a predicate P_2 over the relaxed:

$$\langle P_1 \cdot P_2 \rangle \equiv inj_o(P_1) \wedge inj_r(P_2)$$

Projections. In addition to injection functions, we also define the semantic functions $prj_o(P^*)$ and $prj_r(P^*)$. Each function projects a relational formula in P^* to the set of states corresponding to either the first (prj_o) or second (prj_r) component of each state pair in the denotation of the formula:

$$\begin{aligned}
prj_o(P^*) &\equiv \{\sigma_1 \mid (\sigma_1, \sigma_2) \in \llbracket P^* \rrbracket\} \\
prj_r(P^*) &\equiv \{\sigma_2 \mid (\sigma_1, \sigma_2) \in \llbracket P^* \rrbracket\}
\end{aligned}$$

The projection functions allow us to decompose a relational formula about the behavior of both the original semantics and the relaxed semantics into the set of states that satisfy the relation for either the original or relaxed semantics individually. We can use projection to define the following relations between relational and non-relational formulas:

$$\begin{aligned}
P^* \models_o P &\equiv prj_o(P^*) \subseteq \llbracket P \rrbracket \\
P^* \models_r P &\equiv prj_r(P^*) \subseteq \llbracket P \rrbracket
\end{aligned}$$

Fresh Variables and Substitution. The predicate $fresh(x)$, denoting that x is a fresh variable in the context of an inference rule, is true if $x \in Vars$ and x does not appear in the rule's premises or consequent. Our proof rules also use the standard capture-avoiding substitution $P[e/x]$. We represent multiple substitution $P[e_1/x_1] \cdots [e_n/x_n]$ as $P[e_1, \dots, e_n/x_1, \dots, x_n]$. Substitution over P^* is defined in the same way.

Auxiliary Notations. We also define the following judgments for later use in both the rules of our program logics and the discussion of their semantics:

$$\begin{aligned}
\sigma \models P &\equiv \sigma \in \llbracket P \rrbracket & (\sigma_1, \sigma_2) \models P^* &\equiv (\sigma_1, \sigma_2) \in \llbracket P^* \rrbracket \\
\models P_1 \Rightarrow P_2 &\equiv \llbracket P_1 \rrbracket \subseteq \llbracket P_2 \rrbracket & \models P_1^* \Rightarrow P_2^* &\equiv \llbracket P_1^* \rrbracket \subseteq \llbracket P_2^* \rrbracket
\end{aligned}$$

3.2 Original Semantics

Figure 7 presents a manual translation of our Coq formalization of the axiomatic original semantics of the program. The Hoare-style judgment $\vdash_o \{P\} s \{Q\}$ models the original execution of the program wherein relax statements have no effect. The intended meaning of the judgment is the semantic judgment $\models_o \{P\} s \{Q\}$: for all states σ , if $\sigma \models P$ and $\langle s, \sigma \rangle \Downarrow_o \langle \sigma', \psi \rangle$, then $\sigma' \models Q$. Namely, if σ satisfies P , then if an original execution of s from σ yields a new state σ' , then σ' satisfies Q . We note that this definition asserts only partial correctness and not total correctness.

We have elided a discussion of the rules for standard constructs (i.e., `skip`, `assign`, `sequential composition`, `if`, `while`, and `consequence`) because their definitions are the same as in standard presentations (e.g., Floyd and Hoare [14, 15]). We define the non-standard rules as follows:

The havoc rule requires in its premise that e should be satisfiable by adjusting only the variables in X while retaining the other variables' values. The havoc statement evaluates to wr only if there is no way to transform the current state into a new state satisfying e just by changing the variables in X .

The assert rule asserts that a formula e is true at a given point in the program's execution. This rule therefore requires that e hold in the precondition.

The assume rule assumes that a formula e is true at a given point in the program's execution. The rule differs from the rule for `assert` statements in that it assumes the validity of e and then makes e part of the consequent. Because there is no obligation to prove e for an `assume` statement, e may not hold for all states that satisfy P and, as a result, the `assume` may evaluate to ba . However, by design, we allow `assume` statements to fail in the dynamic original semantics.

The relax rule specifies that a `relax` statement is a no-op and does not change the program's state in the original semantics. However, the current state of the program must also still satisfy e in order to be consistent with our definition of relaxation. The rule reuses the rule for `assert` to enforce this and prevent the dynamic semantics from evaluating to wr .

The relate rule gives `relate` statements the same semantics as `skip` because, unlike the axiomatic relaxed semantics (Section 3.2), the axiomatic original semantics references only a single execution of the program and does not use relational reasoning.

3.3 Relaxed Semantics

Figure 8 presents a manual translation of our Coq formalization of the axiomatic relaxed semantics of the program. The proof rules for the relaxed semantics are relational in that they relate executions of the program under the dynamic relaxed semantics to executions under the dynamic original semantics. The intended meaning of the judgment is the partial correctness assertion $\models_r \{P^*\} s \{Q^*\}$: if $(\sigma_o, \sigma_r) \models P^*$, and $\langle s, \sigma_o \rangle \Downarrow_o \langle \sigma'_o, \psi_1 \rangle$, and $\langle s, \sigma_r \rangle \Downarrow_r \langle \sigma'_r, \psi_2 \rangle$, then $(\sigma'_o, \sigma'_r) \models Q^*$.

The rules are designed to transfer the reasoning done by a developer in the axiomatic original semantics to prove properties about the relaxed semantics. Specifically, the axiomatic relaxed semantics need not re-prove properties about the dynamic original semantics—e.g., the validity of `assert` statements; it can simply

$$\boxed{\vdash_o \{P\} s \{Q\}}$$

$$\begin{array}{c}
\text{skip} \frac{}{\vdash_o \{P\} \text{skip} \{P\}} \quad \text{seq} \frac{\vdash_o \{P\} s_1 \{R\} \quad \vdash_o \{R\} s_2 \{Q\}}{\vdash_o \{P\} s_1; s_2 \{Q\}} \quad \text{assign} \frac{}{\vdash_o \{Q[e/x]\} \mathbf{x} = e \{Q\}} \\
\text{havoc} \frac{\llbracket (\exists X'. P[X'/X]) \wedge e \rrbracket \neq \emptyset \quad \text{fresh}(X')}{\vdash_o \{P\} \text{havoc}(X) \text{st}(e) \{(\exists X'. P[X'/X]) \wedge e\}} \quad \text{assert} \frac{}{\vdash_o \{P \wedge e\} \text{assert } e \{P \wedge e\}} \\
\text{assume} \frac{}{\vdash_o \{P\} \text{assume } e \{P \wedge e\}} \quad \text{relax} \frac{\vdash_o \{P\} \text{assert } e \{Q\}}{\vdash_o \{P\} \text{relax}(X) \text{st}(e) \{Q\}} \quad \text{if} \frac{\vdash_o \{P \wedge b\} s_1 \{Q\} \quad \vdash_o \{P \wedge \neg b\} s_2 \{Q\}}{\vdash_o \{P\} \text{if}(b) \{s_1\} \text{else} \{s_2\} \{Q\}} \\
\text{relate} \frac{}{\vdash_o \{P\} \text{relate } l : e^* \{P\}} \quad \text{while} \frac{\vdash_o \{P \wedge b\} s \{P\}}{\vdash_o \{P\} \text{while}(b) \{s\} \{P \wedge \neg b\}} \quad \text{conseq} \frac{\models P \Rightarrow P' \quad \vdash_o \{P'\} s \{Q'\} \quad \models Q' \Rightarrow Q}{\vdash_o \{P\} s \{Q\}}
\end{array}$$

Figure 7. Axiomatic Original Semantics

$$\boxed{\vdash_r \{P^*\} s \{Q^*\}}$$

$$\begin{array}{c}
\text{relax} \frac{\llbracket (\exists X'(\mathbf{r}). P^*[X'(\mathbf{r})/X(\mathbf{r})]) \wedge \text{inj}_r(e) \rrbracket \neq \emptyset \quad \text{fresh}(X')}{\vdash_r \{P^*\} \text{relax}(X) \text{st}(e) \{(\exists X'(\mathbf{r}). P^*[X'(\mathbf{r})/X(\mathbf{r})]) \wedge \langle e \cdot e \rangle\}} \quad \text{relate} \frac{}{\vdash_r \{P^* \wedge e^*\} \text{relate } l : e^* \{P^* \wedge e^*\}} \\
\text{assert} \frac{\models P^* \wedge \text{inj}_o(e) \Rightarrow \text{inj}_r(e)}{\vdash_r \{P^*\} \text{assert } e \{P^* \wedge \langle e \cdot e \rangle\}} \quad \text{if} \frac{\models P^* \Rightarrow \langle b \cdot b \rangle \vee \langle \neg b \cdot \neg b \rangle \quad \vdash_r \{P^* \wedge \langle b \cdot b \rangle\} s_1 \{Q^*\} \quad \vdash_r \{P^* \wedge \langle \neg b \cdot \neg b \rangle\} s_2 \{Q^*\}}{\vdash_r \{P^*\} \text{if}(b) \{s_1\} \text{else} \{s_2\} \{Q^*\}} \\
\text{assume} \frac{\models P^* \wedge \text{inj}_o(e) \Rightarrow \text{inj}_r(e)}{\vdash_r \{P^*\} \text{assume } e \{P^* \wedge \langle e \cdot e \rangle\}} \quad \text{while} \frac{\models P^* \Rightarrow \langle b \cdot b \rangle \vee \langle \neg b \cdot \neg b \rangle \quad \vdash_r \{P^* \wedge \langle b \cdot b \rangle\} s \{P^*\}}{\vdash_r \{P^*\} \text{while}(b) \{s\} \{P^* \wedge \langle \neg b \cdot \neg b \rangle\}} \\
\text{diverge} \frac{P^* \models_o P_o \quad P^* \models_r P_r \quad \vdash_o \{P_o\} s \{Q_o\} \quad \vdash_r \{P_r\} s \{Q_r\} \quad \text{no_rel}(s)}{\vdash_r \{P^*\} s \{\langle Q_o \cdot Q_r \rangle\}}
\end{array}$$

Figure 8. Axiomatic Relaxed Semantics

assume that these properties are established by the axiomatic original semantics and then transfer their validity to the relaxed semantics via relational reasoning. We have elided discussions of most of the standard rules; however, the non-standard rules are as follows:

The relax rule distinguishes the semantics of relaxation in the original and relaxed semantics of the program. The rule is similar to the havoc rule in the axiomatic original semantics except that 1) it deals with a relational formula and 2) the rule only modifies (i.e., substitutes) relaxed variables, such as $\mathbf{x}(\mathbf{r})$, whereas variables over the original semantics $\mathbf{x}(o)$ are not modified. We use a shorthand $X(\mathbf{r}) \equiv \{x(\mathbf{r}) \mid x \in X\}$ to denote the syntactic extension of a set of plain variables in X to relaxed variables.

The relate rule enables reasoning about `relate` statements in the program. Similar to a non-relational assertion, the rule requires that e^* hold in the precondition. This ensures that e^* holds for all pairs of original and relaxed executions that reach the statement.

The assert rule demonstrates how we can use relational reasoning to prove assertions in the relaxed semantics. Specifically, we can first assume that the assertion is true in the original semantics (i.e., $\text{inj}_o(e)$) because it has been verified with the axiomatic original semantics. We can then conclude that the assertion is true in the relaxed semantics (i.e., $\text{inj}_r(e)$) if $P^* \wedge \text{inj}_o(e)$ implies $\text{inj}_r(e)$.

For example, if the precondition of the statement `assert e` says that all the free variables in e are the same in both semantics (written $\bigwedge_{x \in \text{free}(e)}. \mathbf{x}(o) == \mathbf{x}(\mathbf{r})$), then because the axiomatic original semantics proves that the assertion is true in the dynamic original semantics, we can conclude that the assertion is also true in the dynamic relaxed semantics.

The assume rule demonstrates how relational reasoning allows us to use relations between the original and relaxed semantics of the

program to reason about assumptions in the relaxed semantics in the same way as we do for assertions — i.e., if the assumption is true in the original semantics of the program, then it is also true in the relaxed semantics.

Convergent and Divergent Control Flow. An important aspect of relaxed programs is that the original and relaxed executions of a program may branch in different directions at a control flow construct. Specifically, if two executions branch in the same direction, then we can continue to reason about them relationally in lockstep. However, if the two executions diverge, then the executions execute different statements and, as a result, we lose our relational reasoning power. We note it is possible for two executions to diverge at a control flow construct and then converge again at the end of the construct, allowing us to regain our relational reasoning power.

The relaxed axiomatic semantics captures this property via a set of proof rules for *convergent control flow* constructs (i.e., the original and relaxed executions always branch in the same direction) and another set for *divergent control flow* constructs.¹

3.3.1 Convergent Control Flow

The `if` rule allows us to continue using relational reasoning inside its branches if control flow is convergent. This is established by checking that for all σ_1, σ_2 , if $(\sigma_1, \sigma_2) \models P^*$ then the conditional's boolean expression either evaluates to *true* in both the original and relaxed semantics or it evaluates to *false* in both the original and relaxed semantics. This means that in all cases, the original and relaxed semantics take the same branch together. If this condition is not satisfied, then the rule cannot be applied.

¹We provide a detailed formalization of convergent control flow in the supplementary material distributed with this paper.

$$\boxed{\vdash_i \{P\} s \{Q\}} \quad \text{relax} \frac{\vdash_i \{P\} \text{havoc}(X) \text{st}(e) \{Q\}}{\vdash_i \{P\} \text{relax}(X) \text{st}(e) \{Q\}} \\
\text{assume} \frac{}{\vdash_i \{P \wedge e\} \text{assume } e \{P \wedge e\}}$$

Figure 9. Axiomatic Intermediate Semantics

The `while` rule is similar in form to the `if` rule in that it requires that control flow be convergent to allow us to continue using relational reasoning within the body of a `while` statement.

3.3.2 Divergent Control Flow

The `diverge` rule enables a proof to proceed if the original and relaxed semantics diverge at a control flow construct. Namely, the rule establishes the consequent of the statement by independently establishing that $\vdash_o \{P_o\} s \{Q_o\}$ for the original semantics and that $\vdash_i \{P_r\} s \{Q_r\}$ for the relaxed semantics, where P_o and P_r are left and right projections of P^* . The judgment $\vdash_i \{P_r\} s \{Q_r\}$ is a set of proof rules for the axiomatic *intermediate* semantics of the program. Figure 9 gives an abbreviated presentation of the program logic for the intermediate semantics of the program.

The axiomatic intermediate semantics is a non-relational characterization of the dynamic relaxed semantics and is very similar to the axiomatic original semantics. The intended meaning of the judgment is $\models_i \{P\} s \{Q\}$: for all states σ , if $\sigma \models P$ and $\langle s, \sigma \rangle \Downarrow_r \langle \sigma', \psi \rangle$, then $\sigma' \models Q$. This semantics differs from the axiomatic original semantics in two ways:

- **The relax rule** specifies that `relax` (X) `st` (e) modifies the variables in X such that their values satisfy e whereas in the axiomatic original semantics the `relax` statement is a no-op. To capture this, the rule uses the `havoc` rule to specify the effects.
- **The assume rule** differs from its counterpart in the original semantics because the effects of `relax` statements may violate the programmer’s original assumptions about the behavior of the program. For example, relaxation may interfere with the values of the variables in a subsequent `assume` statement. Therefore, to ensure that an `assume` statement does not evaluate to `ba`, the rule requires that e hold in the precondition — just as in an `assert` statement.

We conclude the presentation of the `diverge` rule by noting that it is also guarded by the predicate $no_rel(s)$, which evaluates to `true` if no `relate` statements appear within s . This predicate therefore prevents `relate` statements from appearing in divergent control flow statements where we are unable to use relational reasoning to establish that the `relate` statement is satisfied.

We also note that the use of projections by the \models_o and \models_r relations in this rule means that all relational properties between the two semantics are lost and must be reestablished at the end of the statement. Relational properties that are not modified by the statement, however, can be preserved via a relational frame rule.

4. Properties

We now present the technical definitions, lemmas, and theorems that establish the semantic properties of programs in the language. In our formalization, we restrict ourselves to terminating relaxed programs. Also, while we only present brief proof sketches, the full sources of our Coq formalization and proofs are available online at

<http://groups.csail.mit.edu/pac/acceptability>

4.1 Original Semantics

Our axiomatic definition for the original semantics is sound and can be used to establish a *weak* form of the traditional progress

theorem for programs. Namely, if you can write a proof in the axiomatic original semantics and an evaluation in the original semantics terminates, as given by the rules, then the resulting state is not *wr*. This differs from a *strong* form of progress that establishes the same for all programs (including non-terminating), which would require a small-step or coinductive formalization of our dynamic semantics.

Lemma 1 (Soundness).

$$\text{If } \vdash_o \{P\} s \{Q\}, \text{ then } \models_o \{P\} s \{Q\}$$

This lemma establishes that our axiomatic definition is sound with respect to the dynamic original semantics of the program. More specifically, given a proof $\vdash_o \{P\} s \{Q\}$, it is the case that for all states $\sigma \models P$, if $\langle s, \sigma \rangle \Downarrow_o \langle \sigma', \psi \rangle$, then $\sigma' \models Q$.

Proof Sketch. This proof proceeds by induction on the rules of $\vdash_o \{P\} s \{Q\}$. A large portion of the proof effort involves proving the semantics of substitution in the case of the assignment rule and the `havoc` rule. The case of the `havoc` rule also requires mutual induction on the lists of modified and fresh variables to establish that the post-condition holds. The cases for structural rules (`if` and sequential composition) follow from induction and the case for the `while` statement proceeds by nested induction on derivations of the evaluation relation. \square

Lemma 2 (Original Progress Modulo Assumptions).

$$\text{If } \vdash_o \{P\} s \{Q\}, \text{ and } \sigma \models P, \text{ and } \langle s, \sigma \rangle \Downarrow_o \phi, \text{ then } \phi \neq wr.$$

This lemma establishes the progress property that we desire for the original semantics. Namely, given a proof in the original axiomatic semantics, then for all states that satisfy P , if evaluation terminates, then the evaluation does not yield *wr*. By design, the judgment does not preclude the program from evaluating to *ba* (indicating that it has violated an assumption).

Proof Sketch. This proof proceeds by induction on the rules of $\vdash_o \{P\} s \{Q\}$. We only need to consider three primitive statements where the program may evaluate to *wr*: `havoc` — the satisfiability check in the premise guards against this; `assert` — the fact that e must hold in the precondition guards against this; and `relax` — follows by induction because evaluation of a `relax` statement reduces to evaluation of an `assert` statement. \square

4.2 Intermediate Semantics

The axiomatic relaxed semantics establishes several progress properties about the relaxed execution of the program. However, to prove progress for the axiomatic relaxed semantics, we need to first prove the same for the axiomatic intermediate semantics.

The definition for the axiomatic intermediate semantics is sound and, also, establishes a form of progress for the relaxed semantics of the program. Specifically, the intermediate semantics models the behavior of a relaxed execution after it has branched at a control flow construct in a different direction than an original execution. When this happens, the relaxed execution must not violate assertions (evaluate to *wr*) or violate assumptions (evaluate to *ba*).

Lemma 3 (Soundness).

$$\text{If } \vdash_i \{P\} s \{Q\}, \text{ then } \models_i \{P\} s \{Q\}$$

This lemma establishes that our axiomatic definition is sound with respect to the dynamic relaxed semantics. More specifically, given a proof $\vdash_i \{P\} s \{Q\}$, it is the case that for all states $\sigma \models P$, if $\langle s, \sigma \rangle \Downarrow_r \langle \sigma', \psi \rangle$, then $\sigma' \models Q$.

Proof Sketch. This proof proceeds by induction on the rules of $\vdash_i \{P\} s \{Q\}$. Because the axiomatic intermediate semantics reuses a large portion of the axiomatic original semantics, and the definition of the dynamic relaxed semantics is very similar to that of the dynamic original semantics, the vast majority of the proof follows from the proofs in Lemma 1. \square

Lemma 4 (Progress).

If $\vdash_i \{P\} s \{Q\}$, **and** $\sigma \models P$, **and** $\langle s, \sigma \rangle \Downarrow_r \phi$,
then $\neg err(\phi)$
where $err(\phi) \equiv \phi = wr \vee \phi = ba$.

This lemma establishes the progress property that we desire for the axiomatic intermediate semantics. Namely, given a proof $\vdash_i \{P\} s \{Q\}$, it is the case that for all states that satisfy P , if evaluation terminates, then the evaluation does not yield an error. Note that this guarantee is stronger than that of the axiomatic original semantics in that it does not allow a relaxed execution to violate an assumption whereas an original execution may do so.

Proof Sketch. The proof proceeds by induction on the rules of $\vdash_i \{P\} s \{Q\}$. As in the proof of soundness (Lemma 3), the vast majority of the proof follows from our proofs about the axiomatic original semantics; in this case most of the proof follows directly from Lemma 2 (Original Progress Modulo Assumptions). The proof differs for two statements: `relax` — the proof follows by induction because evaluation of a `relax` statement reduces to the evaluation of a `havoc`; and `assume` — the proof also follows by induction as evaluation of an `assume` reduces to an `assert`. \square

4.3 Relaxed Semantics

Lemma 5 (Soundness).

If $\vdash_r \{P^*\} s \{Q^*\}$, **then** $\models_r \{P^*\} s \{Q^*\}$.

This lemma establishes that our axiomatic definition is sound with respect to the original and relaxed semantics of the program. Specifically, given a proof $\vdash_r \{P^*\} s \{Q^*\}$, it is the case that for all states $(\sigma_o, \sigma_r) \models P^*$, if $\langle s, \sigma_o \rangle \Downarrow_o \langle \sigma'_o, \psi_1 \rangle$ and $\langle s, \sigma_r \rangle \Downarrow_r \langle \sigma'_r, \psi_2 \rangle$, then $(\sigma'_o, \sigma'_r) \models Q^*$.

Proof Sketch. The proof proceeds by induction on the rules of $\vdash_r \{P^*\} s \{Q^*\}$. The proof is largely similar to that for the original axiomatic semantics in that much of the work lies in proving the semantics of substitution for the `relax` statement, which has a proof that is similar to `havoc` in the original axiomatic semantics. The cases for structural rules (if and sequential composition) follow from induction and the case for the while statement proceeds by nested, mutual induction on derivations of the original and relaxed execution of the statement. The most unique case is the rule for `diverge`: this proof uses the soundness of the original axiomatic semantics (Lemma 1) and the soundness of the intermediate axiomatic semantics (Lemma 3) to establish the soundness of the rule for evaluations in which the original and relaxed executions branch in different directions at a control flow construct. \square

Theorem 6 (Relative Relaxed Progress).

If $\vdash_r \{P^*\} s \{Q^*\}$, **and** $(\sigma_o, \sigma_r) \models P^*$, **and** $\langle s, \sigma_o \rangle \Downarrow_o \phi_o$,
and $\neg err(\phi_o)$, **and** $\langle s, \sigma_r \rangle \Downarrow_r \phi_r$, **then** $\neg err(\phi_r)$

This theorem establishes the relative progress guarantee for the relaxed semantics of the program. Namely, given a proof $\vdash_r \{P^*\} s \{Q^*\}$, it is the case that for all pairs of states (σ_o, σ_r) that satisfy P^* , if an original execution terminates and does not produce an error, then if a relaxed execution terminates, it also does not produce an error.

Proof Sketch. This proof proceeds by induction on the rules of $\vdash_r \{P^*\} s \{Q^*\}$. The most important cases are the `assert` and `assume` statements and the `diverge` rule. The proofs for `assert` and `assume` are similar in that the premise ensures that if the original execution evaluates to *true*, then the condition also evaluates to *true* in the relaxed execution. We do not have to consider the case where the original execution evaluates to *false* because it would mean that $\phi_o = wr \vee \phi_o = ba$, which is inconsistent with our theorem's assumption that $\neg err(\phi_o)$.

The `diverge` rule demonstrates the utility of our design of the axiomatic intermediate semantics. For this rule, we can no longer use facts about the original execution to prove facts about the relaxed execution. Therefore, the relaxed execution must be inherently error free. The proof uses the progress guarantee of the axiomatic intermediate semantics (Lemma 4) to establish exactly that. \square

Theorem 7 (Relaxed Progress).

If $\vdash_o \{P\} s \{Q\}$, **and** $\vdash_r \{P^*\} s \{Q^*\}$, **and** $P^* \models_o P_o$,
and $(\sigma_o, \sigma_r) \models P^*$, **and** $\langle s, \sigma_o \rangle \Downarrow_o \phi_o$, **and** $\phi_o \neq ba$,
and $\langle s, \sigma_r \rangle \Downarrow_r \phi_r$, **then** $\neg err(\phi_r)$.

This theorem combines the multiple proofs and assumptions in our programming model to establish the final progress guarantee for relaxed programs: given 1) a proof in the original axiomatic semantics, 2) a proof in the relaxed axiomatic semantics, and 3) that executions in the original semantics terminate and do not violate an assumption, then if the relaxed program terminates, it does not produce an error.

Proof Sketch. This proof follows directly from our assumptions, Lemma 2 (Original Progress Modulo Assumptions), and Theorem 6 (Relative Progress). Namely, our assumptions and Lemma 2 establish that if an original execution terminates, then it does not terminate in error. By Theorem 6, we can conclude that if a relaxed execution terminates, it does not produce an error. \square

Corollary 8 (Relaxed Progress Modulo Original Assumptions).

If $\vdash_o \{P\} s \{Q\}$, **and** $\vdash_r \{P^*\} s \{Q^*\}$, **and** $P^* \models_o P_o$,
and $(\sigma_o, \sigma_r) \models P^*$, **and** $\langle s, \sigma_o \rangle \Downarrow_o \phi_o$, **and** $\langle s, \sigma_r \rangle \Downarrow_r \phi_r$,
and $err(\phi_r)$, **then** $\phi_o = ba$

This variant of Theorem 7 captures an important aspect of the design of our programming model in that our language enables developers to discern the provenance of errors. Namely, this corollary states precisely that given a proof in the original axiomatic semantics and a proof in the relaxed axiomatic semantics, if the original and relaxed evaluations of the program terminate, then if an error occurs in the relaxed evaluations, then it is due to an original evaluation that violates an assumption.

Theorem 9 (Relational Soundness).

If $\vdash_r \{P^*\} s \{Q^*\}$, **and** $(\sigma_o, \sigma_r) \models P^*$,
and $\langle s, \sigma_o \rangle \Downarrow_o \langle \sigma'_o, \psi_1 \rangle$, **and** $\langle s, \sigma_r \rangle \Downarrow_r \langle \sigma'_r, \psi_2 \rangle$,
then $\Gamma \vdash \psi_1 \sim \psi_2$

This theorem establishes that given a proof in the relaxed axiomatic semantics, if the original evaluation of the program terminates successfully and the relaxed evaluation of the program terminates successfully, then the observation lists generated by the executions $(\psi_1$ and ψ_2 , respectively) satisfy the observational compatibility relation $\Gamma \vdash \psi_1 \sim \psi_2$. Observational compatibility implies that the original and relaxed executions of the program satisfy all executed `relate` statements.

Figure 10 presents the definition of the observational compatibility relation. The symbol Γ represents a finite map from `relate`

$$\boxed{\Gamma \vdash \psi_1 \sim \psi_2} \quad \overline{\Gamma \vdash \emptyset \sim \emptyset}$$

$$\frac{l_1 = l_2 \quad \llbracket \Gamma(l_1) \rrbracket(\sigma_1, \sigma_2) = \text{true} \quad \Gamma \vdash \psi_1 \sim \psi_2}{\Gamma \vdash (l_1, \sigma_1) :: \psi_1 \sim (l_2, \sigma_2) :: \psi_2}$$

Figure 10. Observational Compatibility Relation

labels to relational boolean expressions (i.e., $\Gamma \in L \rightarrow B^*$). We define this map by structural induction on the syntax of the program, where the label of each `relate` statement in the program maps to its relational boolean expression. We require that `relate` statements in well-formed programs be uniquely labeled.

The rules specify that if two observations lists are empty, then they are compatible. Otherwise, any two lists are compatible if 1) the labels in the head are the same (indicating that they are generated by the same `relate` statement), 2) the relational boolean expression for the label evaluates to `true` for the states in the head, and 3) the tails of the two lists are also compatible.

Proof Sketch. This proof proceeds by induction on the rules of $\vdash_r \{P^*\} s \{Q^*\}$. The two interesting cases are the *diverge* rule and the rule for `relate` statements. For the *diverge* rule, we use the fact that the rule requires *no_rel*(*s*) hold, which means that no `relate` statements appear inside *s*. We can therefore conclude that the observation list for the statement is empty and, therefore, original and relaxed executions of the statement are trivially compatible. In the case of the `relate` rule, the proof uses the rule’s precondition to establish that the two emitted observations satisfy the `relate` statement’s condition. \square

5. Example Relaxed Programs

Inspired by programs that researchers have successfully relaxed in prior work, we developed several example programs designed to capture the core aspects of the successful relaxations. We then formalized key acceptability properties of the relaxations and used our Coq formalization to prove these properties.

5.1 Dynamic Knobs

Swish++ is an open-source search engine. We work with a successful relaxation that uses Dynamic Knobs to reduce the number of search results that Swish++ presents to the user when the server is under heavy load [16]. The rationale for this relaxation is that 1) users are typically only interested in the top search results and 2) users are very sensitive to how quickly the results are presented — even a short delay can significantly reduce advertisement revenue.

Relaxation. The transformation targets a loop that formats and presents the search query results. The loop keeps track of the number of search results, which we denote by *N*. The loop also has a control variable `max_r` which is a threshold on the number of elements that should be presented to the user: if *N* is smaller than `max_r`, then all results will be presented; otherwise, only the first `max_r` results will be presented.

A relaxed program can nondeterministically change `max_r` to reduce the number of iterations of this loop while still returning the most important results:

```

original_max_r = max_r;
relax (max_r) st
  (original_max_r <= 10 && max_r == original_max_r)
  || (10 < original_max_r && 10 <= max_r);

```

This code first saves the original value of the control variable `max_r` in `original_max_r`. It then relaxes `max_r`. There are two cases: if the original value of this control variable was less than

or equal to 10, then the relaxed execution should be the same as the original execution — it presents the same number of results, since the value of `max_r` does not change. If, on the other hand, the original value was greater than 10, the only constraint is that the value of `max_r` is not smaller than 10, meaning that it should return at least the top 10 results when available. The `relax` statement nondeterministically changes `max_r` subject to these constraints.

Acceptability. One acceptability property is that the relaxed execution must present either all of the search results from the original execution to the user (if the number of search results in the original execution is less than or equal to 10), or at least the first 10 results (if the number of results in the original execution is greater than 10). The following `relate` statement captures these constraints:

```

relate (num_r<o> < 10 && num_r<o> == num_r<r>) ||
  (10 <= num_r<o> && 10 <= num_r<r>);

```

The loop that formats and presents the search results maintains a count `num_r` of the number of formatted and presented results. This statement therefore uses the value of `num_r` in the original program (denoted `num_r<o>`) to determine how many search results the original execution presents. The `relate` statement uses `num_r<o>` and the (potentially different) value of `num_r` in the relaxed execution (`num_r<r>`) to formalize the desired relationship between the two executions.

Verification. The proof of the `relate` statement involves 330 lines of Coq proof scripts. Because the relaxation changes the number of loop iterations, the proof uses the divergent control flow rule to reason about the loop in the original semantics and relaxed semantics separately. The key proof steps establish that the condition of the `relax` statement holds before entering the loop and that `original_max_r<o> == original_max_r<r>` and `N<o> = N<r>`. The loop invariant in both the original and relaxed execution is `num_r <= max_r && num_r <= N`.

Once control flow converges after the loop, the `relate` statement’s precondition can be deduced via a proof by cases or, as in our proof environment, verified by an automated theorem prover.

5.2 Statistical Automatic Parallelization

Our next example is drawn from a parallelization of the Water computation [8] with statistical accuracy bounds [20]. In this computation a control variable determines whether to execute a loop sequentially or in parallel. To maximize performance, the parallelization eliminates lock operations that make updates to an array *RS*. The resulting race conditions produce a parallel computation whose result may vary nondeterministically (because of processor scheduling variations) within acceptable accuracy bounds.

Relaxation. We model the relaxation nondeterminism by relaxing each element in *RS* with no constraints:²

```

relax (RS) st (true);

```

In a loop that executes after parallelization, the Water computation compares *RS*[*K*] to a cutoff variable `gCUT2` and, if it is less than the cutoff, uses *RS*[*K*] to update an array *FF* (here `EXP(RS[K])`) is an expression involving *RS*):

```

while (K < N) {
  if (RS[K] < gCUT2) { FF[K] = EXP(RS[K]); }
  K = K + 1;
}

```

² Although we do not present our treatment of arrays in this paper, our Coq formalization supports reasoning about arrays, which is a straightforward extension of our presented framework.

Acceptability. A key acceptability property is that K stays within the bounds of the array `FF`.³ The array bounds are stored in the variable `len_FF`. We assume that the developer establishes, via some standard reasoning process, that the original execution does not violate the array bounds and therefore inserts the statement `assume (K < len_FF)`, inside the `if` statement, just before the assignment to `FF[K]`.

Verification. Recall that the verification of the relaxed program must verify that the condition in each `assume` statement holds in the relaxed execution. One approach is noninterference — verify that relaxation does not affect the values of the variables in the predicate. However, this is a relational property and because the `assume` statement appears at a divergent control flow point (it is affected by the value of the relaxed variable `RS`), this approach does not work.

The developer therefore inserts another `assume` statement, `assume (K < len_FF)`, just before the `if` statement. It is possible to verify this statement using noninterference and then propagate the condition through the `if` statement to verify the second `assume` statement.

The Coq verification of this program consists of approximately 310 lines of proof script. The key prerequisite of the proof is to establish the relational loop invariants that $K\langle o \rangle == K\langle r \rangle$ and $\text{len_FF}\langle o \rangle == \text{len_FF}\langle r \rangle$. These invariants enable us to prove that the relaxation does not interfere with the assumption.

5.3 Approximate Memory and Data Types

Our third example is drawn from the LU decomposition algorithm implemented in the SciMark2 benchmark suite [2]. Researchers have demonstrated that lower-power, approximate memories and CPU compute units can be used to lower the energy consumption of this computation at the expense of a small loss in accuracy [18, 33].

We focus on the part of the computation that computes the pivot row `p` for each column `j` in a matrix `A`. The pivot row is the row that contains the maximum element in the column.

```
i = j + 1;
while ( i < N ) {
  a = A[i][j];
  if ( a > max ) { max = a; p = i; }
  i = i + 1;
}
```

Relaxation. Following the assumptions on errors in approximate memories described in [25], if `A` is stored in approximate memory, then we can model the range of errors when reading a value from `A` with a relaxation that non-deterministically adds bounded error (ϵ) to the result:

```
original_a = a;
relax (a) st (original_a - e <= a &&
  a <= original_a + e);
```

Acceptability. One acceptability property for this computation is that the value in the selected pivot row (`max`) in the relaxed execution does not differ from the result in an original execution by more than ϵ . We can specify this with a `relate` statement:

```
relate max<o> - max<r> <= e && max<r> - max<o> <= e
```

We note that this `relate` statement asserts the *Lipschitz-continuity* of the computation: small changes in the inputs lead to proportionally small changes in the output.

³ We note that K must also be within the bounds of `RS`; the proof is similar.

Verification. The Coq verification of this program consists of approximately 315 lines of proof script. The key prerequisite of the proof is to establish that the relational property in the given `relate` statement is loop invariant.

6. Future Work

Termination. Our proof rules are designed to support the verification of safety properties, which must hold for all executions of a program, and relational properties, which relate original and relaxed executions. While these properties establish a partial correctness guarantee as in standard Hoare logic, some relaxed programs may also require total correctness: proofs of termination as well as partial correctness.

Proving a relaxed program’s termination can be more challenging than the traditional termination problem because relaxed programs can be extended with additional nondeterminism that may change the control flow dynamically and nondeterministically. One possible way to produce termination proofs for a relaxed program is to reuse the termination proof of the original program in the same way we have done for safety properties. While our proof machinery is not designed to support termination proofs, we believe that the rules in the axiomatic relaxed semantics establish a relative termination guarantee for convergent while loops — loops with conditions that can be proved to be equivalent in both the original and relaxed semantics. Relative termination for these loops would mean that if the loop terminates in the original program then it also terminates in the relaxed program. We anticipate that such a notion of relative termination (akin to our relative progress) will be a fruitful direction for future work.

Completeness. We have omitted a formal statement of the completeness of the logic because we believe that there are additional interesting properties of relaxed programs that interact with termination, divergent control flow, and the semantics of errant executions in which the original or relaxed semantics can evaluate to *ba* or *wr*. We intend to explore the subject of completeness and its interactions with these properties in future work.

7. Related Work

Executable Specifications. Executable specifications, via techniques such as refinement and constraint solving, produce concrete outputs that satisfy the specification [13, 19, 23, 27, 32, 36]. Applications include recovering from errors in existing code and providing alternate implementations for code that may be difficult to develop using standard techniques.

The research in this paper differs in that it promotes nondeterministic relaxation to obtain semantically different but still acceptable variants of the original program. Our focus is therefore on enabling developers to specify and prove correctness requirements as relational properties between the original and relaxed program.

Unreliable Memory and Critical Data. Researchers have proposed techniques for enabling programs to distinguish data that can be stored in unreliable low-power memory from critical data whose values must be reliably stored [10, 18, 33]. These systems typically focus on data values (such as the values of pixels in an image) that can, in principle, legally take on *any* value. The techniques presented in this paper make it possible to prove general correctness properties of these kinds of programs (whose relaxations may primarily involve unreliably stored data that can take on any value).

Relational Program Logics. Our program logic for the relaxed semantics of the program builds on previous work on the Relational Hoare Logic (RHL) [7]. RHL itself was inspired by work on Translation Validation [26] and Credible Compilation [31] and has since inspired other forms of relational reasoning. Researchers have

also defined relational separation logic [5, 35], probabilistic Hoare logic [6], and have used relational reasoning to verify the correctness of semantics-preserving loop optimizations [12], Lipschitz-continuity [12], access control policies [24], and differential privacy mechanisms [6].

While majority of the previous research has been focused on proving that transformed programs retain the semantics of the original version, our goal is different — specifically, to prove that relaxed executions (which typically have different semantics) preserve important correctness properties. We adapt RHL to prove properties that relate the original and relaxed executions and extend RHL to reason about assertions (which reference only the current execution) and assumptions (which are assumed to hold in original executions but must be shown to hold in relaxed executions).

8. Conclusion

The additional nondeterminism in relaxed programs enables programs to operate at a variety of points with different combinations of accuracy, performance, and resource consumption characteristics. It is possible to exploit this flexibility to satisfy a variety of goals, including trading off accuracy for enhanced performance or reduced energy consumption [3, 4, 11, 16–18, 20–22, 30, 33, 34, 37] or responding to load spikes or other fluctuations in the characteristics of the underlying computational platform [16, 17, 29, 33].

We present formal reasoning techniques that make it possible to verify important correctness properties of relaxed programs. Standard verification techniques reference only the current execution of the current program under verification. Our techniques, in contrast, aim to reduce the verification effort by taking a relational approach that exploits the close relationship between the original and relaxed executions. Our goal is to give developers the verified correctness properties they need to confidently deploy relaxed programs and exploit the substantial flexibility, performance, and resource consumption advantages that relaxed programs offer.

Acknowledgments

We would like to thank Fan Long and our anonymous reviewers for their insightful comments. We also note our earlier technical report on reasoning about relaxed programs [9]. This research was supported in part by the National Science Foundation (Grants CCF-0811397, CCF-0905244, CCF-1036241, and IIS-0835652), DARPA (Grants FA8650-11-C-7192 and FA8750-12-2-0110), and the United States Department of Energy (Grant DE-SC0005288).

References

- [1] The Coq Proof Assistant. <http://coq.inria.fr>.
- [2] Scimark 2.0. <http://math.nist.gov/scimark2>.
- [3] J. Ansel, C. Chan, Y. L. Wong, M. Olszewski, Q. Zhao, A. Edelman, and S. Amarasinghe. Petabricks: a language and compiler for algorithmic choice. PLDI, 2009.
- [4] W. Baek and T. M. Chilimbi. Green: a framework for supporting energy-conscious programming using controlled approximation. PLDI, 2010.
- [5] G. Barthe, J. Crespo, and C. Kunz. Relational verification using product programs. FM, 2011.
- [6] G. Barthe, B. Köpf, F. Olmedo, and S. Zanella Béguelin. Probabilistic reasoning for differential privacy. POPL, 2012.
- [7] N. Benton. Simple relational correctness proofs for static analyses and program transformations. POPL, 2004.
- [8] W. Blume and R. Eigenmann. Performance analysis of parallelizing compilers on the Perfect Benchmarks programs. *Transactions on Parallel and Distributed Systems*, 3(6), 1992.
- [9] M. Carbin, D. Kim, S. Misailovic, and M. Rinard. Reasoning about Relaxed Programs. Technical Report MIT-CSAIL-TR-2011-050, MIT, 2011.
- [10] M. Carbin and M. Rinard. Automatically Identifying Critical Input Regions and Code in Applications. ISSSTA, 2010.
- [11] S. Chaudhuri, S. Gulwani, R. Lubliner, and S. Navidpour. Proving Programs Robust. FSE, 2011.
- [12] J.M. Crespo and C. Kunz. A machine-checked framework for relational separation logic. SEFM, 2011.
- [13] B. Demsky and M. Rinard. Data structure repair using goal-directed reasoning. ICSE, 2005.
- [14] R. W. Floyd. Assigning meanings to programs. *Mathematical aspects of computer science*, 19(19-32), 1967.
- [15] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10), October 1969.
- [16] H. Hoffman, S. Sidiroglou, M. Carbin, S. Misailovic, A. Agarwal, and M. Rinard. Dynamic knobs for responsive power-aware computing. ASPLOS, 2011.
- [17] H. Hoffmann, S. Misailovic, S. Sidiroglou, A. Agarwal, and M. Rinard. Using Code Perforation to Improve Performance, Reduce Energy Consumption, and Respond to Failures. Technical Report MIT-CSAIL-TR-2009-042, MIT, 2009.
- [18] S. Liu, K. Pattabiraman, T. Moscibroda, and B. Zorn. Flikker: Saving dram refresh-power through critical data partitioning. ASPLOS, 2011.
- [19] A. Milicevic, D. Rayside, K. Yessenov, and D. Jackson. Unifying execution of imperative and declarative code. ICSE, 2011.
- [20] S. Misailovic, D. Kim, and M. Rinard. Parallelizing sequential programs with statistical accuracy tests. Technical Report MIT-CSAIL-TR-2010-038, MIT, 2010.
- [21] S. Misailovic, D. Roy, and M. Rinard. Probabilistically Accurate Program Transformations. SAS, 2011.
- [22] S. Misailovic, S. Sidiroglou, H. Hoffmann, and M. Rinard. Quality of service profiling. ICSE, 2010.
- [23] C. Morgan. The specification statement. *Transactions on Programming Languages and Systems*, 10(3), 1988.
- [24] A. Nanevski, A. Banerjee, and D. Garg. Verification of information flow and access control policies with dependent types. SP, 2011.
- [25] J. Nelson, A. Sampson, and L. Ceze. Dense approximate storage in phase-change memory. ASPLOS-WACI, 2011.
- [26] A. Pnueli, M. Siegel, and E. Singerman. Translation validation. TACAS, 1998.
- [27] D. Rayside, A. Milicevic, K. Yessenov, G. Dennis, and D. Jackson. Agile specifications. OOPSLA, 2009.
- [28] M. Rinard. Acceptability-oriented computing. OOPSLA Onwards '03.
- [29] M. Rinard. Probabilistic accuracy bounds for fault-tolerant computations that discard tasks. ICS, 2006.
- [30] M. Rinard. Using early phase termination to eliminate load imbalances at barrier synchronization points. OOPSLA, 2007.
- [31] M. C. Rinard and D. Marinov. Credible compilation with pointers. RTRV, 1999.
- [32] H. Samimi, E. Aung, and T. Millstein. Falling back on executable specifications. ECOOP, 2010.
- [33] A. Sampson, W. Dietl, E. Fortuna, D. Gnanapragasam, L. Ceze, and D. Grossman. Enerj: approximate data types for safe and general low-power computation. PLDI, 2011.
- [34] S. Sidiroglou, S. Misailovic, H. Hoffmann, and M. Rinard. Managing Performance vs. Accuracy Trade-offs With Loop Perforation. FSE '11.
- [35] H. Yang. Relational separation logic. *Theoretical Computer Science*, 375(1-3), May 2007.
- [36] J. Yang, K. Yessenov, and A. Solar-Lezama. A language for automatically enforcing privacy policies. POPL, 2012.
- [37] Z. Zhu, S. Misailovic, J. Kelner, and M. Rinard. Randomized accuracy-aware program transformations for efficient approximate computations. POPL, 2012.