RHLE: Modular Deductive Verification of Relational ∀∃ Properties*

Robert Dickerson, Qianchuan Ye, Michael K. Zhang, and Benjamin Delaware

Purdue University, West Lafayette, IN 47907, USA
rob@robd.io, ye202@purdue.edu,
michael.k.zhang@alumni.purdue.edu, bendy@purdue.edu

Abstract. Hoare-style program logics are a popular and effective technique for software verification. Relational program logics are an instance of this approach that enable reasoning about relationships between the execution of two or more programs. Existing relational program logics have focused on verifying that *all* runs of a collection of programs do not violate a specified relational behavior. Several important relational properties, including refinement and noninterference, do not fit into this category, as they also mandate the *existence* of specific desirable executions. This paper presents RHLE, a logic for verifying these sorts of relational ∀∃ properties. Key to our approach is a novel form of function specification that employs a variant of ghost variables to ensure that valid implementations exhibit certain behaviors. We have used a program verifier based on RHLE to verify a diverse set of relational ∀∃ properties drawn from the literature.

1 Introduction

Hoare-style program logics are a popular and effective verification technique. Starting with Hoare's seminal paper [18], this approach has been adapted to cover a variety of programming languages and assertions [3, 19, 30, 32, 26]. These logics typically feature several pleasant properties: they can be declaratively specified via a set of rules over the syntax of the target programming language, they permit compositional reasoning over individual program components, and they often admit effective automated verification procedures. Most of these logics focus on proving *safety* properties of *single* programs, i.e., that executing a program in a valid initial state never results in a state violating a postcondition.

Not all program behaviors fall into this category, however. As one example, consider the common scenario where a developer decides they want to migrate a hand-rolled implementation of a function to one that uses a third-party library. Figure 1 gives a concrete example of this situation. The program on the left, $\texttt{sample}_1$, uses a random number generator to directly sample a subset of an array. The program on the right, $\texttt{sample}_2$, opts to delegate the task to an external list library which supports shuffling and constructing sublists. While $\texttt{sample}_1$ works *with replacement* (the same elements may be sampled multiple times), $\texttt{sample}_2$

* This research was partially supported by the National Science Foundation under Grant CCF-1755880 and by a grant from the Purdue Research Foundation.

```
int[] sample_1(int[] arr,
               int size) {
  assert(size <= arr.length);
  int[] samp = new int[size];
  for (i in [0..size]) {
    int j = randB(arr.length);
    samp[i] = arr[j];
  }
  return samp;
}
```

```
int[] sample_2(int[] arr,
               int size) {
  assert(size <= arr.length);
  list = new List(arr);
  perm = list.permute();
  samp = perm.sublist(size);
  return samp.toArray();
}
```

Fig. 1. An example migration of a function which randomly samples a list of integers with replacement to a function which samples without replacement. The original program (sample$_1$) uses a function which generates random numbers, while the migrated program (sample$_2$) uses a list abstraction with a permute operation.

works *without replacement* (an element may be sampled at most once). In order to ensure that this change does not break things, the developer may wish to verify that sample$_2$ does not do anything that sample$_1$ could not, i.e., that the updated function *refines* the original. Notably, this refinement property relates the behavior of *multiple* programs. In addition, it does not have the form of a standard safety property. The developer does not want to enforce that sample$_2$ produces *every* permutation that the hand-rolled implementation does; rather, they wish to ensure it does not start returning previously impossible samples.

```
int encode(int msg^H) {
  int key^H = randB(MAX_INT);
  int enc^L = msg^H xor key^H;
  return (key^H, enc^L);
}
```

As another example, consider the encode function on the right which performs a simple xor cipher. This function takes a single high-security argument, msg$^{\mathrm{H}}$, and returns a pair of high-security and low-security results, key$^{\mathrm{H}}$ and enc$^{\mathrm{L}}$, respectively. The function encodes its argument by first generating a random key (randB returns a random value between 0 and its argument), taking the xor of the key and the message, and finally returning the key along with the encoded message. The developer may wish to guarantee an attacker can learn nothing about the secret message given only the encoded message. Whether or not encode meets this *generalized noninterference* [24] property crucially depends on the behavior of randB: if the attacker knows this function *always* returns 3, for example, they can decipher any encoded message. We can again frame this behavior as a relational property between the executions of two programs (in this case calls to encode with arbitrary arguments msg$^{\mathrm{H}}_1$ and msg$^{\mathrm{H}}_2$): every execution of encode(msg$^{\mathrm{H}}_1$) must have a corresponding execution of encode(msg$^{\mathrm{H}}_2$) that returns the same low-security encoded value.

In both examples, the desired behavior has the shape *for all* executions of some program, *there exists* a corresponding execution of a second program that is somehow related. Thus, we call these properties *relational* $\forall\exists$ properties. While several *relational program logics* have been developed for reasoning about the behavior of multiple programs [34, 9, 8], all have focused on relational *safety* properties, i.e., that *all* the final states of multiple programs satisfy some

$$
\begin{array}{ll}
n \in \mathbb{N} \qquad x, y \in \mathcal{V} & s ::= \mathsf{skip} \mid s;\ s \\
f, g \in \mathcal{N} \qquad \sigma \in \mathcal{V} \to \mathbb{N} & \quad \mid \mathsf{if}\ b\ \mathsf{then}\ s\ \mathsf{else}\ s \\
a ::=\ n \mid x \mid a + a \mid a - a \mid a * a & \quad \mid \mathsf{while}\ b\ \mathsf{do}\ s\ \mathsf{end} \\
b ::=\ \mathsf{true} \mid \mathsf{false} & \quad \mid x := a \mid x := \mathsf{havoc} \mid x := f(\overline{a}) \\
\quad \mid a = a \mid a < a \mid \neg b \mid b \wedge b & FD ::=\ \textbf{def}\ f(\overline{x})\ \{s; \textbf{return}\ a\}
\end{array}
$$

Fig. 2. Syntax of FunIMP.

relational postcondition. Unfortunately, in the presence of nondeterminism, none of these logics are capable of verifying relational ∀∃ properties such as refinement and generalized noninterference. The need to reason about nondeterminism naturally arises in the presence of external functions like `permute` in Figure 1, where specifications are used to approximate the behavior of multiple possible implementations.

This paper addresses this gap by introducing RHLE, a relational program logic for reasoning about ∀∃ properties. Key to our approach is a novel form of function specifications which approximate the set of behaviors a valid implementation *must* exhibit. These specifications use a novel variant of ghost variables, which we call *choice variables*, which guarantee the existence of required behaviors. RHLE admits a modular reasoning principle, where any properties verified against a set of function specifications continue to hold whenever the program is linked to any satisfying implementation. While techniques based on Constrained Horn Clauses [36] and model checking [23] have recently been developed that are capable of reasoning about ∀∃ properties, RHLE is, to the best of our knowledge, the first Hoare-style program logic for doing so. We have used a verifier based on RHLE to verify a range of ∀∃ properties including refinement, noninterference (with and without delimited release), semantic parameter usage, and flaky tests.

We begin by defining a core imperative language with function calls (Section 2) equipped with semantics for both over- and under-approximating function behaviors (Section 3). We next present RHLE, and a corresponding verification algorithm for verifying ∀∃ properties (Section 5). We evaluate our approach by applying an implementation of this algorithm to verify a diverse set of relational properties (Section 6). We conclude with an examination of related work (Section 7). We have formalized the details of our approach in the Coq proof assistant; this development is available in the supplementary materials of this paper.

2 The FunIMP Language

We begin with the definition of FunIMP, a core imperative language with function calls $x := f(\overline{a})$ and nondeterministic variable assignment $x := \mathsf{havoc}$. The full syntax of FunIMP is presented in Figure 2. The calculus is parameterized over disjoint sets of identifiers for program variables $\mathcal{V}$ and function names $\mathcal{N}$. Functions have a fixed arity. Function definitions consist of a sequence of statements followed by an expression that computes the result of the function. For brevity, we denote sequences $x_1, \ldots, x_n$ as $\overline{x}$. For ease of presentation, we treat functions as returning a single value, although it is straightforward to extend

FunIMP to allow for multiple return values: $(x, y, \ldots) := f(\overline{a})$. Our verification tool, ORHLE (see Section 6), uses such an extension to model functions which mutate their arguments.

The semantics of FunIMP programs are defined via a standard big-step evaluation relation from initial to final program states. States are mappings from variables to integers, and are usually notated as σ. We write $[x \mapsto a]\sigma$ to refer to state σ updated with a mapping from x to a. The evaluation rules are parameterized over an *implementation context*, a mapping $I \in \mathcal{N} \to FD$ from function names to their definitions, which is used to evaluate function calls:

$$\frac{\begin{array}{c} I(f) = \textsf{\textbf{def}}\ f(\overline{x})\,\{s; \textsf{\textbf{return}}\ e\} \\ I \vdash \sigma, \overline{a} \Downarrow \overline{v} \qquad I \vdash [\overline{x} \mapsto \overline{v}], s \Downarrow \sigma' \qquad I \vdash \sigma', e \Downarrow r \end{array}}{I \vdash \sigma, y := f(\overline{a}) \Downarrow [y \mapsto r]\sigma}\ \textsc{ECall}$$

We use $\Downarrow$ for the evaluation relation of both expressions and statements; $\sigma, e \Downarrow \sigma'$ holds when executing e on state σ can result in state σ'. Since programs may be nondeterministic, there may be multiple final states related to a single initial state for a given program. Note that $\textsf{havoc}$ is the only source of nondeterminism when evaluating a FunIMP program. The remaining evaluation rules for FunIMP are standard and can be found in Appendix A.

3 Approximating FunIMP Behaviors

In order to modularly reason about relational $\forall\exists$ properties, we first present semantics for capturing the possible executions of a FunIMP program in *any* valid implementation context. In order to account for both “for all” and “there exists” behaviors of functions, we rely on two kinds of specifications. To reason about *all* possible executions of a valid implementation, i.e., a standard *safety* property, we use a *universal* specification. For guarantees about the *existence* of certain executions, we use an *existential* specification.

3.1 Universal Executions

Both kinds of specifications are parameterized over an assertion language $\mathcal{A}$ on program states and a mechanism for judging when a state satisfies an assertion. We write $\sigma \models P$ to denote that a state σ satisfies the assertion P. The universal specifications used to reason about programs on the “for all” side of $\forall\exists$ properties are written as $FA ::= \textsf{ax}_\forall\ f(\overline{x})\ \{P\}\{Q\}$, where $P \in \mathcal{A}$ is a precondition with free variables in $\overline{x}$ and $Q \in \mathcal{A}$ is a postcondition with free variables in $\overline{x} \cup \{\rho\}$. The postcondition uses the distinguished variable ρ to refer to the value returned by f. Universal specifications promise client programs that the valid implementations of a function will only evaluate to states satisfying the postcondition when evaluated in a starting state that satisfies the precondition.

Definition 1 ($\forall - Compatibility$)**.** *A function definition* $\textsf{def}\ f(\overline{x})\{s; \textsf{return}\ r\}$ *is $\forall$-compatible with a universal specification* $\textsf{ax}_\forall\ f(\overline{x})\{P\}\{Q\}$ *if only values* *satisfying Q may be returned whenever f is called with arguments satisfying P:*

$$\forall \sigma, \sigma'.\ (\sigma \models P)\ \wedge\ (I \vdash \sigma, s \Downarrow \sigma')\ \wedge\ (\sigma', r \Downarrow v)\ \Longrightarrow\ ([\rho \mapsto v]\sigma \models Q)$$

We say that an implementation context I is ∀-compatible with a context of universal specifications $S_\forall \in \mathcal{N} \rightarrow FA$ when every definition in I is ∀-compatible with the corresponding specification in $S_\forall$.

To characterize the set of possible behaviors of a program under any ∀-compatible implementation context, we define a new *overapproximate* semantics for FunIMP, $\Downarrow_\forall$. The evaluation rules of this semantics are based on $\Downarrow$, but they use a universal specification context, $S_\forall$, instead of an implementation context, and replace ECall with the following two evaluation rules:

$$\frac{\begin{array}{c} S_\forall(f) = \mathsf{ax}_\forall\ f(\overline{x})\ \{P\}\ \{Q\} \\ S_\forall \vdash \sigma, \overline{a} \Downarrow_\forall \overline{v} \qquad [\overline{x} \mapsto \overline{v}] \models P \qquad [\rho \mapsto r, \overline{x} \mapsto \overline{v}] \models Q \end{array}}{S_\forall \vdash \sigma, y := f(\overline{a}) \Downarrow_\forall [y \mapsto r]\sigma}\ \textsc{ECall}_{\forall 1}$$

$$\frac{S_\forall(f) = \mathsf{ax}_\forall\ f(\overline{x})\ \{P\}\ \{Q\} \qquad S_\forall \vdash \sigma, \overline{a} \Downarrow_\forall \overline{v} \qquad [\overline{x} \mapsto \overline{v}] \not\models P}{S_\forall \vdash \sigma, y := f(\overline{a}) \Downarrow_\forall [y \mapsto r]\sigma}\ \textsc{ECall}_{\forall 2}$$

The first rule states that If a function is called with arguments satisfying its precondition, it will return a value satisfying its postcondition; otherwise, the second rule states that it can return *any* value. The latter case allows the overapproximate semantics to capture evaluations where a function is called with arguments that do not meet its precondition. Appendix A includes a complete listing of the $\Downarrow_\forall$ relation.

Any final state of a program evaluated under an implementation context I which is ∀-compatible with $S_\forall$ can also be produced using $\Downarrow_\forall$ and $S_\forall$. Appealing to this intuition, we call the evaluations of a FunIMP program p using $\Downarrow_\forall$ the *overapproximate executions* of p under $S_\forall$.

Theorem 1. *When run under an implementation context I that is ∀-compatible with specification context $S_\forall$ and an initial state σ, a program p will either diverge or evaluate to a state σ' which is also the result of one of its overapproximate executions under $S_\forall$.*

3.2 Existential Executions

Universal specifications approximate function calls on the “for all” side of ∀∃ properties by constraining what a compatible implementation *can* do. Existential specifications approximate the “there exists” executions by describing the required values a valid implementation *must* be able to return. In order to flexibly capture these behaviors, existential pre- and post-conditions are indexed by a set of *choice variables* $\overline{c} \subseteq \mathcal{V}$. Each instantiation of these variables defines a particular behavior that an implementation has to exhibit. The syntax for writing an existential specification is: $FE ::= \mathsf{ax}_\exists\ f(\overline{x})\ [\overline{c}]\ \{P\}\{Q\}$.

```
def randB(x) {
  skip;
  return 0
}
```

```
def randB(x) {
  r := havoc;
  while (x ≤ r) do r := r − x end;
  return r }
```

```
def randB(x) {
  r := havoc;
  return r
}
```

Fig. 3. Implementations of a function which returns an integer within a bound.

We write $A[x/y]$ to denote the predicate A with all free occurrences of x replaced with y. Intuitively, for any instantiation $\overline{v}$ of choice variables $\overline{c}$, an existential specification requires an implementation to produce at least one value satisfying the specialized postcondition $Q[\overline{v}/\overline{c}]$, when called with arguments that satisfy the corresponding precondition $P[\overline{v}/\overline{c}]$. This intuition is embodied in our notion of compatibility for existential specifications:

Definition 2 (∃-Compatibility). *A function definition* $\mathsf{def}\ f(\overline{x})\{s; \mathsf{return}\ r\}$ *is ∃-compatible with an existential specification* $\mathsf{ax}_\exists\ f(\overline{x})[\overline{c}]\{P\}\{Q\}$ *if, for every selection of choice variables* $\overline{v}$*, calling* f *with arguments that satisfy* $P[\overline{v}/\overline{c}]$ *can return at least one value satisfying* $Q[\overline{v}/\overline{c}]$*:*

$$\forall \sigma, \overline{v}.\ (\sigma \models P[\overline{v}/\overline{c}]) \implies \exists \sigma'.\ (I \vdash \sigma, s \Downarrow \sigma')\ \wedge\ (\sigma', r \Downarrow v) \wedge ([\rho \mapsto v]\sigma \models Q[\overline{v}/\overline{c}])$$

Example 1. To see how universal and existential specifications work together to describe a function's behavior, consider a function randB(x) which is intended to return some integer between 0 and its argument x. We can write a universal specification requiring all return values to be within the desired bound: $\mathsf{ax}_\forall$ $\mathsf{randB}(x)$ $\{0 < x\}$ $\{0 \leq \rho < x\}$. This does not, however, *guarantee* every value in this range is possible. To express this requirement, we reify the choice of the random value using an existential specification: $\mathsf{ax}_\exists$ $\mathsf{randB}(x)$ $[c]$ $\{0 < x \wedge 0 \leq c < x\}$ $\{\rho = c\}$. Figure 3 lists a variety of possible randB implementations; the first implementation is compatible with the aforementioned universal specification and the third definition is compatible with the existential specification, but only the middle one satisfies both. Note how c acts as a ghost variable which constrains the choice of the random number. Thus, when reasoning about a client of randB, we can select a concrete value for c that forces the desired result.

Equipped with a context of existential specifications $S_\exists \in \mathcal{N} \rightarrow FE$, we characterize the set of behaviors a program *must* exhibit under every ∃-compatible implementation context via an underapproximate semantics for FunIMP programs. The judgements of this semantics are denoted as $S_\exists\ \vdash \sigma, p \Downarrow_\exists \Sigma$, which reads as: under context $S_\exists$ and initial state σ, the program p will produce at least one final state in the set of states Σ. The evaluation rules of this semantics are given in Figure 4. Most of the rules in Figure 4 adapt the FunIMP evaluation rules to account for the fact that commands now produce *sets* of states from an initial state. For example, the evaluation rule for sequences, ESeq$_\exists$, states that s_2 produces a final state corresponding to every state in the set produced by s_1. The rule for function calls, ECall$_\exists$, is the most interesting: it *chooses* one of the behaviors guaranteed by the existential specification of a function and produces a *set* of final states for every return value consistent with that choice.

$$\frac{}{S_\exists \vdash \sigma, \mathsf{skip} \Downarrow_\exists \{\sigma\}}\ \textsc{ESkip}_\exists \qquad \frac{}{S_\exists \vdash \sigma, x := \mathsf{havoc} \Downarrow_\exists \{\sigma' \mid \exists v.[x \mapsto v]\sigma'\}}\ \textsc{EHavoc}_\exists$$

$$\frac{\sigma, a \Downarrow v}{S_\exists \vdash \sigma, x := a \Downarrow_\exists \{[x \mapsto v]\sigma\}}\ \textsc{EAssn}_\exists \qquad \frac{S_\exists \vdash \sigma, s \Downarrow_\exists \Sigma \qquad \Sigma \subseteq \Sigma'}{S_\exists \vdash \sigma, s \Downarrow_\exists \Sigma'}\ \textsc{EConsq}_\exists$$

$$\frac{S_\exists \vdash \sigma, s_1 \Downarrow_\exists \Sigma \qquad \forall \sigma' \in \Sigma.\ S_\exists \vdash \sigma', s_2 \Downarrow_\exists \Sigma'}{S_\exists \vdash \sigma, s_1;\ s_2 \Downarrow_\exists \Sigma'}\ \textsc{ESeq}_\exists$$

$$\frac{\begin{array}{c}\sigma, b \Downarrow \mathsf{true} \qquad S_\exists \vdash \sigma, c \Downarrow_\exists \Sigma \\ \forall \sigma' \in \Sigma.\ S_\exists \vdash \sigma', \mathsf{while}\ b\ \mathsf{do}\ c\ \mathsf{end} \Downarrow_\exists \Sigma'\end{array}}{S_\exists \vdash \sigma, \mathsf{while}\ b\ \mathsf{do}\ c\ \mathsf{end} \Downarrow_\exists \Sigma'}\ \textsc{ELpT}_\exists$$

$$\frac{\sigma, b \Downarrow \mathsf{false}}{S_\exists \vdash \sigma, \mathsf{while}\ b\ \mathsf{do}\ c\ \mathsf{end} \Downarrow_\exists \{\sigma\}}\ \textsc{ELpF}_\exists \qquad \frac{\sigma, b \Downarrow \mathsf{true} \qquad S_\exists \vdash \sigma, s_1 \Downarrow_\exists \Sigma}{S_\exists \vdash \sigma, \mathsf{if}\ b\ \mathsf{then}\ s_1\ \mathsf{else}\ s_2 \Downarrow_\exists \Sigma}\ \textsc{EIfT}_\exists$$

$$\frac{\sigma, b \Downarrow \bot \qquad S_\exists \vdash \sigma, s_2 \Downarrow_\exists \Sigma}{S_\exists \vdash \sigma, \mathsf{if}\ b\ \mathsf{then}\ s_1\ \mathsf{else}\ s_2 \Downarrow_\exists \Sigma}\ \textsc{EIfF}_\exists$$

$$\frac{S_\exists(f) = \mathsf{ax}_\exists\ f(\overline{x})\,[c]\,\{P\}\,\{Q\} \qquad S_\exists \vdash \sigma, \overline{a} \Downarrow \overline{v} \qquad [\overline{x} \mapsto \overline{v}] \models P[\overline{k}/\overline{c}]}{S_\exists \vdash \sigma, y := f(\overline{a}) \Downarrow_\exists \quad \{\sigma' \mid \exists r.\ \sigma' = [y \mapsto r]\sigma\ \wedge\ [\rho \mapsto r, \overline{x} \mapsto \overline{v}] \models Q[\overline{k}/\overline{c}]\}}\ \textsc{ECall}_\exists$$

Fig. 4. The existential evaluation relation.

Every set of final states for a program p produced by these semantics under $S_\exists$ includes a possible final state of p when evaluated under any ∃-compatible implementation context. For this reason, we term the evaluations of p using $\Downarrow_\exists$ the *underapproximate executions* of p under $S_\exists$.

Theorem 2. *If there is an underapproximate evaluation of program p to a set of states Σ from an initial state σ under $S_\exists$, then p must terminate in at least one final state $\sigma' \in \Sigma$ when it is run from σ under an implementation context I that is ∃-compatible with $S_\exists$.*

3.3 Approximating ∀∃ behaviors

Taken together, the over- and under-approximate semantics allow us to relate the ∀∃ behaviors of multiple client programs under every ∀- and ∃-compatible implementation context. This admits a modular reasoning principle, where if a set of clients can be shown to exhibit some behaviors using the overapproximate and underapproximate semantics, linking the client with any compatible environment will continue to exhibit those behaviors. The key challenge to ensuring these ∀∃ behaviors is identifying, for every overapproximate execution, an appropriate selection of choice variables that cause the underapproximate executions to evaluate to a collection of final states satisfying a desired ∀∃ property.

Example 2. Consider the second example from the introduction, and assume that `randB` has the universal and existential specifications from Example 1. To ensure that `encode` does not reveal anything about its secret input via its public output, it suffices to establish that for any universal execution of `encode` on a specific input, every other possible input to `encode` could produce the same encoded message under the existential semantics. The first execution begins with the statement **int** $\texttt{key}^{\texttt{H}}_{\forall}$`= randB(MAX_INT)` (for convenience, we annotate program variables from the first and second executions with the subscripts ∀ and ∃, respectively). By $\text{E}\textsc{Call}_{\forall 1}$, this statement will update $\texttt{key}^{\texttt{H}}_{\forall}$ to hold a value between 0 and `MAX_INT`. The function then encodes the message using this key, and returns the result. In order to show this leaks nothing, we need to establish a corresponding execution of `encode` that returns this same result regardless of the value of its argument. In effect, this amounts to finding a strategy for instantiating the choice variable in $\text{E}\textsc{Call}_{\exists}$ to assign an appropriate value to $\texttt{key}^{\texttt{H}}_{\exists}$. In this case, the choice is straightforward: we need a c such that c `xor` $\texttt{msg}^{\texttt{H}}_{\exists}$ $=\texttt{enc}^{\texttt{L}}_{\forall}$. Using $\texttt{msg}^{\texttt{H}}_{\exists}$ `xor` $\texttt{enc}^{\texttt{L}}_{\forall}$ for c in $\text{E}\textsc{Call}_{\exists}$ achieves the desired result. Using this strategy, we can construct an appropriate execution in response to *every* execution of `encode`. In contrast, if our existential specification were $\mathsf{ax}_{\exists}$ $\mathsf{randB}(x)$ $[\,]$ $\{0 < x\}$ $\{0 \leq \rho < x\}$, it would only guarantee the existence of a single result, and there would be no workable strategy. Indeed, the first definition of `randB` in Figure 3 satisfies this specification, and `encode` will always leak the full message when using this implementation!

4 RHLE

We now present RHLE, a relational program logic for proving that a collection of FunIMP programs exhibit some desired set of ∀∃ behaviors. As a consequence of Theorem 1 and Theorem 2, this entails that properties established in RHLE will continue to hold when the programs are linked with any compatible implementation context.

RHLE specifications use *relational* assertions (denoted $\Phi, \Psi \in \mathcal{A}$) to relate the execution of multiple programs. As normal assertions are predicates on a single state, a relational assertion is a predicate on multiple states. Each program in a RHLE triple operates over a distinct state space. To disambiguate between variables that occur in multiple copies, shared variable names are annotated with an identifier unique to each program. Following existing convention [34, 9], we use a natural number to identifying which state a variable belongs to. As an example, the relational assertion $x_1 \leq x_2$ is a binary predicate over (at least) two states. This assertion is satisfied by any set of two (or more) states where the value of x in the first state is less than or equal to the value of x in the second.

RHLE triples have the form $S_\forall, S_\exists \models \langle \Phi \rangle\, \overline{p_\forall} \sim_{\exists} \overline{p_\exists}\, \langle \Psi \rangle$ and assert that *for all* universal executions of the programs $\overline{p_\forall}$, *there exist* existential executions of the programs $\overline{p_\exists}$ satisfying the relational pre- and post-condition Φ and Ψ:

$$
\begin{aligned}
S_\forall, S_\exists \models \langle \Phi \rangle\, \overline{p_\forall} \sim_{\exists} \overline{p_\exists}\, \langle \Psi \rangle \equiv \quad & \forall \overline{\sigma_\forall}\; \overline{\sigma_\exists}\; \overline{\sigma'_\forall}.\;\; \overline{\sigma_\forall},\; \overline{\sigma_\exists} \models \Phi \;\wedge\; S_\forall \vdash \overline{\sigma_\forall},\; \overline{p_\forall} \Downarrow_\forall \overline{\sigma'_\forall} \implies \\
& \exists \Sigma.\; S_\exists \vdash \overline{\sigma_\exists},\; \overline{p_\exists} \Downarrow_\exists \Sigma \;\wedge\; \forall \sigma'_\exists \in \Sigma.\; \overline{\sigma'_\forall},\; \overline{\sigma'_\exists} \models \Psi
\end{aligned}
$$

Property	**RHLE Assertion**
Refinement	$S_\forall,\ S_\exists \models \langle \overline{x_1} = \overline{x_2} \rangle\, y_1 := f(\overline{x_1}) \sim_\exists y_2 := f(\overline{x_2})\, \langle y_1 = y_2 \rangle$
Noninterference	$S_\forall,\ S_\exists \models \langle low_1 = low_2 \rangle\, p_1 \sim_\exists p_2\, \langle low_1 = low_2 \rangle$
Injectivity	$S_\forall,\ S_\exists \models \langle x_1 \neq x_2 \rangle\, y_1 := f(x_1) \circledast y_2 := f(x_2) \sim_\exists \mathsf{skip}\, \langle y_1 \neq y_2 \rangle$
Nondeterminism	$S_\forall,\ S_\exists \models \langle x_1 = x_2 \rangle\, \mathsf{skip} \sim_\exists y_1 := f(x_1) \circledast y_2 := f(x_2)\, \langle y_1 \neq y_2 \rangle$

Table 1. Example RHLE assertions. In the second row, low_x refers to the low security state in program p_x; note the ∀∃ relationship must hold for *any* pair of initial high security values, so $high_x$ is not constrained in the precondition.

We use $\circledast$ to delineate different programs on the universal and existential sides of $\sim_\exists$ so that, e.g., a sequence of n programs $\overline{p}$ is also denoted as $p_1 \circledast \ldots \circledast p_n$. For example, to assert the program $x := \mathsf{havoc}$ is nondeterministic, we write a RHLE triple with two copies of the program, adding a subscript to the variable x in each for clarity: $\cdot \models \langle \top \rangle\, \mathsf{skip} \sim_\exists x_1 := \mathsf{havoc} \circledast x_2 := \mathsf{havoc}\, \langle x_1 \neq x_2 \rangle$. This triple says that, for all starting states and all executions of the trivial program skip, there exist executions of the programs $x_1 := \mathsf{havoc}$ and $x_2 := \mathsf{havoc}$ such that $x_1 \neq x_2$ after both programs have executed. Note that $\circledast$ is *not* a concatenation operator; it does nothing more than delineate multiple programs in a RHLE triple. Table 1 gives some additional examples of RHLE assertions.

$$\frac{}{S_\forall, S_\exists \vdash \langle \Phi \rangle\ \overline{\text{skip}} \sim_\exists \overline{\text{skip}}\ \langle \Phi \rangle}\ \textsc{Finish} \qquad \frac{S_\forall, S_\exists \vdash \langle \Phi \rangle\ \overline{p_\forall;\ \mathsf{skip}} \sim_\exists \overline{p_\exists;\ \mathsf{skip}}\ \langle \Psi \rangle}{S_\forall, S_\exists \vdash \langle \Phi \rangle\ \overline{p_\forall} \sim_\exists \overline{p_\exists}\ \langle \Psi \rangle}\ \textsc{SkipI}$$

$$\frac{\begin{array}{c}\forall \overline{\sigma}\ \overline{\sigma_\exists}.\ S_\forall \vdash \{\Phi \,|_i\, \overline{\sigma},\ \overline{\sigma_\exists}\}\ \ s_i\ \ \{\Phi' \,|_i\, \overline{\sigma},\ \overline{\sigma_\exists}\} \\ S_\forall, S_\exists \vdash \langle \Phi' \rangle\ p_1 \circledast \ldots \circledast s_i' \circledast \ldots \circledast s_n \sim_\exists \overline{p_\exists}\ \langle \Psi \rangle\end{array}}{S_\forall, S_\exists \vdash \langle \Phi \rangle\ p_1 \circledast \ldots \circledast s_i;\ s_i' \circledast \ldots \circledast p_n \sim_\exists \overline{p_\exists}\ \langle \Psi \rangle}\ \textsc{Step}\forall$$

$$\frac{\begin{array}{c}\forall \overline{\sigma_\forall}\ \overline{\sigma}.\ S_\exists \vdash [\Phi \,|_i\, \overline{\sigma_\forall},\ \overline{\sigma}]\ \ s_i\ \ \left[\Phi' \,|_i\, \overline{\sigma_\forall},\ \overline{\sigma}\right]_\exists \\ S_\forall, S_\exists \vdash \langle \Phi' \rangle\ \overline{p_\forall} \sim_\exists p_1 \circledast \ldots \circledast s_i' \circledast \ldots \circledast p_n\ \langle \Psi \rangle\end{array}}{S_\forall, S_\exists \vdash \langle \Phi \rangle\ \overline{p_\forall} \sim_\exists p_1 \circledast \ldots \circledast s_i;\ s_i' \circledast \ldots \circledast p_n\ \langle \Psi \rangle}\ \textsc{Step}\exists$$

Fig. 5. Core RHLE proof rules.

The core logic of RHLE is given in Figure 5. Relational proofs are built by reasoning about the topmost statement of either one of the universally quantified programs via the Step∀ rule or one of the existentially quantified programs using the Step∃ rule. Once all program statements have been considered, final proof obligations can be discharged using the Finish rule. The SkipI rule is used to ensure that all programs end with skip, so that Finish can be applied. Both Step rules rely on non-relational logics for reasoning about the universal $S_\forall \vdash \{P\}\ p\ \{Q\}$ and existential $S_\exists \vdash [P]\ p\ [Q]_\exists$ behaviors of single statements; we will present the details of both logics shortly. The Step rules employ a projection operation, $\overline{\sigma}|_i\Psi$, which maps a relational assertion to a non-relational one. Given a collection of n states, $\Psi \,|_i\, \overline{\sigma}$ is satisfied by any state σ' which satisfies

Ψ when inserted at the ith position:

$$\sigma' \models \Psi \mid_i \overline{\sigma} \quad \equiv \quad \sigma_1, \ldots, \sigma_{i-1}, \sigma', \sigma_{i+1}, \ldots, \sigma_n \models \Psi$$

In effect, this operation ensures the states of the other programs remain unchanged when reasoning about the ith program in the triple.

Universal Hoare Logic The program logic for universal executions has a standard partial correctness semantics:

$$S_\forall \models \{P\}\ p\ \{Q\} \quad \equiv \quad \forall \sigma, \sigma'.\ \sigma \models P \wedge S_\forall \vdash \sigma, p \Downarrow_\forall \sigma' \implies \sigma' \models Q$$

The rules of this logic are largely standard[1], except for the rule for function calls, which uses a context of universal function specifications:

$$\frac{S_\forall(f) = \mathsf{ax}_\forall\ f(\overline{x})\{P\}\{Q\}}{S_\forall \vdash \left\{\begin{matrix} P[\overline{a}/\overline{x}] \quad \wedge & \\ \forall v. Q[v/\rho; \overline{a}/\overline{x}] & \implies \ R[v/y] \end{matrix}\right\}\ y := f(\overline{a})\ \{R\}}\ \forall\textsc{Spec}$$

Existential Hoare Logic The assertions of our program logic for existential executions say that, for any state meeting the precondition, there *exists* an execution of the program ending in a set of states meeting the post-condition:

$$S_\exists \models [P]\ p\ [Q]_\exists \quad \equiv \quad \forall \sigma.\ \sigma \models P \implies \exists \Sigma.\ S_\exists \vdash \sigma, p \Downarrow_\exists \Sigma \quad \wedge \quad \forall \sigma' \in \Sigma.\ \sigma' \models Q$$

These rules are largely standard *total* Hoare logic rules[2], augmented with a rule for calls to existentially specified functions:

$$\frac{S_\exists(f) = \mathsf{ax}_\exists f(\overline{x})\ [\overline{c}]\ \{P\}\ \{Q\}}{S_\exists \vdash \left[\begin{matrix} \exists \overline{k}.\ ([\overline{x} \mapsto \overline{a}] \models P[\overline{k}/\overline{c}] \\ \wedge \quad \exists v.[\rho \mapsto v, \overline{x} \mapsto \overline{a}] \models Q[\overline{k}/\overline{c}] \\ \wedge \quad \forall v.[\rho \mapsto v, \overline{x} \mapsto \overline{a}] \models Q[\overline{k}/\overline{c}] \\ \implies R[v/y]) \end{matrix}\right]\ y := f(\overline{a})\ [R]_\exists}\ \exists\textsc{Spec}$$

The precondition of this rule is quantified over instantiations $\overline{k}$ of the specification's choice variables. The first of the three conjuncts under this quantifier ensures that the statement is executed in a state satisfying the function's precondition. The next conjunct ensures that the function's post-condition is inhabited. The final conjunct requires that every possible return value satisfying the function's post-condition also satisfies the triple's post-condition.

[1] Appendix B gives a full listing of the rules of this logic.

[2] The full existential logic is presented in Appendix C.

Example 3. Given the existential specification $\mathsf{ax}_\exists\ \mathsf{zeroOrOne}()\ [c]\ \{c = 0 \vee c = 1\}$ $\{\rho = c\}$, we can use ∃Spec (along with the rule for while loops given in Appendix C) to prove the existential assertion $S_\exists \vdash [k = 0]\ \mathsf{while}\ k < 4\ \mathsf{do}\ k := k + \mathsf{zeroOrOne}()\ \mathsf{end}\ \ [k = 4]_\exists$. This loop *could* loop forever by choosing to add 0 to k at every iteration. Nevertheless, by using measure $4 - k$ with the well-founded relation $<$ and instantiating the choice variable with 1 at each iteration, we can prove a terminating path through the program exists.

4.1 Synchronous Rules

While the rules in Figure 5 are sufficient to reason about relational properties, it is possible to lessen the verification burden for structurally similar programs by employing *synchronous rules* which exploit structural similarities between the programs being verified [25]. Reasoning over similar control flow structures in lockstep can reduce the space of states verification must consider and simplify loop invariants. This is particularly useful when reasoning about *hyperproperties* [12], or relational properties on multiple executions of the *same* program. In order to more easily reason about structurally similar programs, RHLE also includes synchronous rules inspired by the Cartesian loop logic presented by Sousa and Dillig [34]. Appendix D includes a full listing of these rules.

Example 4. Consider proving that $\mathsf{while}\ \ (\mathsf{x} < 10)\ \ \mathsf{do}\ \ \mathsf{y}\ \ :=\ \ \mathsf{y} + \mathsf{randB}(9)\ \ \mathsf{end}$ refines $\mathsf{while}\ \ (\mathsf{x} < 10)\ \ \mathsf{do}\ \ \mathsf{y}\ \ :=\ \ \mathsf{y} + \mathsf{randB}(5);\ \mathsf{y}\ \ :=\ \ \mathsf{y} + \mathsf{randB}(6)\ \ \mathsf{end}$. Intuitively, the first program refines the second because the bodies of the loops are themselves refinements. A proof using only the rules in Figure 5 is unable to take advantage of this intuition, however. Instead, the proof requires a sufficiently strong invariant characterizing the behavior of the entire loop on the left, and then an invariant for the righthand program that accounts for the behavior of individual iterations of the lefthand loop.

The SyncLoops rule is designed for this situation:

$$\frac{\begin{array}{c} S_\forall, S_\exists \vdash \langle \mathbb{I} \wedge \bigwedge\limits_{0 \le i \le n} b_i \rangle\ s_0 \circledast \cdots \circledast s_k \sim_\exists s_{k+1} \circledast \cdots \circledast s_n\ \langle \mathbb{I} \rangle \\ \mathbb{I} \wedge \bigwedge\limits_{0 \le i \le n} \neg b_i \implies \Psi \qquad \mathbb{I} \wedge \neg \bigwedge\limits_{0 \le i \le n} b_i \implies \bigwedge\limits_{0 \le i \le n} \neg b_i \end{array}}{\begin{array}{c} S_\forall, S_\exists \vdash \langle \mathbb{I} \rangle\ \mathsf{while}\ b_0\ \mathsf{do}\ s_0\ \mathsf{end} \circledast \cdots \circledast \mathsf{while}\ b_k\ \mathsf{do}\ s_k\ \mathsf{end} \\ \sim_\exists \mathsf{while}\ b_{k+1}\ \mathsf{do}\ s_{k+1}\ \mathsf{end} \circledast \cdots \circledast \mathsf{while}\ b_n\ \mathsf{do}\ s_n\ \mathsf{end}\ \langle \Psi \rangle \end{array}}\ \textsc{SyncLoops}$$

The first premise of this rule says that executing all loop bodies preserves some invariant $\mathbb{I}$, the second ensures the invariant is strong enough to imply the postcondition, and the third requires all loops to end on the same iteration. Since this invariant is reestablished after the execution of every loop body; the invariant that y_1 and y_2 are equal at each iteration suffices to verify this example.

4.2 Soundness

The combination of the core and synchronous rules provide a sound methodology for reasoning about ∀∃ properties:

Theorem 3 (RHLE is Sound). *Suppose $S_\forall, S_\exists \vdash \langle \Phi \rangle\ \overline{p_\forall} \sim_\exists \overline{p_\exists}\ \langle \Psi \rangle$. Then, for any function context I compatible with $S_\forall$ and $S_\exists$, any set of initial states $\overline{\sigma_\forall}$ and $\overline{\sigma_\exists}$ satisfying Φ, and every collection of final states $\overline{\sigma'_\forall}$ of $\overline{p_\forall}$, there must exist a collection of final states produced by $\overline{p_\exists}$ that, together with $\overline{\sigma'_\forall}$, satisfies the relational post-condition Ψ.*

5 Verification

```
Algorithm 1: RHLEVerify
  Inputs : $\Phi$, relational precondition
           $p_\forall$, universal programs
           $p_\exists$, existential programs
           $\Psi$, relational postcondition
  Output : $\langle\Phi\rangle\ p_\forall \sim_\exists p_\exists\ \langle\Psi\rangle$ validity
  begin
      $\overline{\Psi} \leftarrow (\varnothing, \varnothing, \Psi)$
      $(\overline{a}, \overline{e}, \Psi') \leftarrow$
        VCGen ($\overline{\mathit{skip}; p_\forall}, \overline{\mathit{skip}; p_\exists}, \overline{\Psi}$)
      return
        Verify ($\forall \overline{a} \exists \overline{e}.\ \Phi \implies \Psi'$)
```

We now turn to the relational verification algorithm based on RHLE, presented in Algorithm 1. The algorithm is implicitly parameterized over a pair of universal and existential contexts, and `Verify`, a decision procedure for checking validity of a formula in the underlying assertion logic. The bulk of the work is delegated to `VCGen`, presented in Algorithm 2, which builds a weakest relational precondition for the input RHLE triple. The algorithm then checks that the RHLE triple's precondition entails the calculated weakest precondition.

The body of `VCGen` builds a formula by recursively generating verification conditions for the input programs statement by statement. This loop tries to maximize opportunities to apply synchronous rules at each step, as these rules allow us to simultaneously generate proof obligations for multiple subprograms, as discussed in Section 4.1. After establishing there are still program statements to step over (lines 3–4), `VCGen` looks for and processes any trailing program statements which are not loops (lines 5–8), as such statements are not subject to synchronous rule applications. To process individual program statements, `VCGen` relies on a pair of verification condition generators, $\mathtt{VC}_\forall$ and $\mathtt{VC}_\exists$, for the non-relational program logics. These functions are largely standard weakest precondition generators extended with support for existential function calls. The consequents of ∀Spec and ∃Spec immediately yield weakest precondition rules, so that if $S_\forall(f) = \mathsf{ax}_\forall\ f(\overline{x})\{P\}\{Q\}$ and $S_\exists(f) = \mathsf{ax}_\exists f(\overline{x})\ [\overline{c}]\ \{P\}\ \{Q\}$, then:

$$\mathtt{VC}_\forall(\Psi, y := f(\overline{a})) = P[\overline{a}/\overline{x}] \wedge \forall v. Q[v/\rho; \overline{a}/\overline{x}] \implies \Psi[v/y]$$

$$\begin{aligned}\mathtt{VC}_\exists(\Psi, y := f(\overline{a})) = \exists \overline{k}.\ ([\overline{x} \mapsto \overline{a}] \models P[\overline{k}/\overline{c}]\ &\wedge\ \exists v.[\rho \mapsto v, \overline{x} \mapsto \overline{a}] \models Q[\overline{k}/\overline{c}] \\ &\wedge\ \forall v.[\rho \mapsto v, \overline{x} \mapsto \overline{a}] \models Q[\overline{k}/\overline{c}] \implies \Psi[v/y])\end{aligned}$$

If the first three cases fail, the final statements of all the remaining programs are loops. In this case, `VCGen` attempts to simultaneously process the loops (lines 9–19) à la the SyncLoops rule in Example 4. To be eligible for fusion, loops must execute in lockstep. This condition is checked (line 16) before returning; if

Algorithm 2: VCGen

```
   Inputs : $p_\forall$, a set of universal programs
            $p_\exists$, a set of existential programs
            $\overline{\Psi} = (Q_\forall, Q_\exists, \Psi)$, $\Psi$ a postcondition with quantified variables $Q_\forall$, $Q_\exists$
   Output : $(\{v_0, \ldots, v_n\}, \{w_0, \ldots, w_n\}, \Phi)$ such that $v_i$, $w_i$ free in $\Phi$ and
            $\langle \Phi \rangle\ p_\forall \sim_\exists p_\exists\ \langle \Psi \rangle$ is valid if $\forall v_0, \ldots, v_n\ \exists w_0, \ldots w_n.\ \Phi \implies \Psi$
1  begin
2      match $p_\forall \sim_\exists p_\exists$:
3          case $\overline{\mathit{skip}} \sim_\exists \overline{\mathit{skip}}$ do
4              return $\overline{\Psi}$
5          case $\overline{p'_\forall} \circledast (s_1; s_2) \circledast \overline{p''_\forall} \sim_\exists p_\exists$ where $s_2$ not a loop do
6              VCGen($\overline{p'_\forall} \circledast s_1 \circledast \overline{p''_\forall}, p_\exists, VC_\forall(s_2, \overline{\Psi})$)
7          case $p_\forall \sim_\exists \overline{p'_\exists} \circledast (s_1; s_2) \circledast \overline{p''_\exists}$ where $s_2$ not a loop do
8              VCGen($p_\forall, \overline{p'_\exists} \circledast s_1 \circledast \overline{p''_\exists}, VC_\exists(s_2, \overline{\Psi})$)
9          case $\overline{p'_\forall} \circledast s_1; \mathit{if}\ b\ \mathit{then}\ s_t\ \mathit{else}\ s_e \circledast \overline{p''_\forall} \sim_\exists p'_\exists$ do
10             $(Q_\forall, Q_\exists, \Psi_T) \leftarrow$ VCGen($\overline{p'_\forall} \circledast s_1; s_t \circledast \overline{p''_\forall}, p_\exists, b \implies \Psi$))
11             $(Q'_\forall, Q'_\exists, \Psi_E) \leftarrow$ VCGen($\overline{p'_\forall} \circledast s_1; s_e \circledast \overline{p''_\forall}, p_\exists, \neg b \implies \Psi$))
12             return $(Q_\forall \cup Q'_\forall, Q_\exists \cup Q_\exists, \Psi_T \wedge \Psi_E)$
13         case $p_0; \mathit{while}\ b_0\ \mathit{do}\ s_0 \circledast \cdots \circledast p_{i-1}; \mathit{while}\ b_{i-1}\ \mathit{do}\ s_{i-1} \sim_\exists$
            $p'_i; \mathit{while}\ b_i\ \mathit{do}\ s_i \circledast \cdots \circledast p'_n; \mathit{while}\ b_n\ \mathit{do}\ s_n$ do
14             $\mathbb{I} \leftarrow$ FindInvariant($\mathit{while}\ b_0\ \mathit{do}\ s_0 \circledast \cdots \circledast \mathit{while}\ b_{i-1}\ \mathit{do}\ s_{i-1} \sim_\exists$
                 $\mathit{while}\ b_i\ \mathit{do}\ s_i \circledast \cdots \circledast \mathit{while}\ b_n\ \mathit{do}\ s_n$)
15             $(Q'_\forall, Q'_\exists, \Psi_{body}) \leftarrow$ VCGen($s_0 \circledast \cdots \circledast s_{i-1} \sim_\exists s_i \circledast \cdots \circledast s_n, \mathbb{I}$)
16             $\mathit{inductive} \leftarrow \mathbb{I} \wedge \bigwedge_{0 \leq i \leq n} b_i \implies \Psi_{body}$
17             $\mathit{lockstep} \leftarrow \mathbb{I} \wedge \neg \bigwedge_{0 \leq i \leq n} b_i \implies \bigwedge_{0 \leq i \leq n} \neg b_i$
18             $\mathit{post} \leftarrow \mathbb{I} \wedge \bigwedge_{0 \leq i \leq n} \neg b_i \implies \Psi$
19             $(Q_\forall, Q_\exists, \Psi) \leftarrow \overline{\Psi}$
20             if Verify($Q_\forall \cup Q'_\forall,\ Q_\exists \cup Q'_\exists,\ \mathit{inductive} \wedge \mathit{lockstep} \wedge \mathit{post}$) then
21                 VCGen($\overline{p}, \overline{p'}, (Q_\forall, Q_\exists, \mathbb{I})$)
22             else
23                 next case
24         case $\overline{p'_\forall} \circledast (s_1; s_2) \circledast \overline{p''_\forall} \sim_\exists p_\exists$ do
25             VCGen($\overline{p'_\forall} \circledast s_1 \circledast \overline{p''_\forall}, p_\exists, vc_\forall(s_2, \overline{\Psi})$)
26         case $p_\forall \sim_\exists \overline{p'_\exists} \circledast (s_1; s_2) \circledast \overline{p''_\exists}$ do
27             VCGen($p_\forall, \overline{p'_\exists} \circledast s_1 \circledast \overline{p''_\exists}, vc_\exists(s_2, \overline{\Psi})$)
```

loops may execute different numbers of times, the algorithm proceeds to the next match case. If no synchronized reasoning is possible, `VCGen` defaults to stepping over an arbitrary loop in one of the programs (lines 20–23).

`VCGen` is parameterized over a procedure called `FindInvariant`, which acts as an oracle for relational loop invariants. Our prototype implementation of Algorithm 1 currently requires requires loops to be annotated with their invariants; these annotations are used to implement `FindInvariant`. We have experimented with adapting both purely logical [17, 16] and data-driven approaches [28, 29] for invariant inference, but have yet to discover one that is effective for our larger benchmarks. Unlike traditional loop invariants, which must be re-established on every possible execution of the loop body, invariants in existentially quantified executions need only be re-established on a subset of the possible execution of the body. A robust invariant inference approach thus requires finding not only the invariant itself, but a strategy for instantiating choice variables that consistently re-establish the chosen invariant. Scalable invariant inference for existentially quantified executions is an important and interesting future direction.

Appendix F includes an example application of Algorithm 1 to `RandB`.

6 Implementation and Evaluation

To evaluate our approach, we have implemented ORHLE, an automatic program verifier based on Algorithm 1. ORHLE is implemented in Haskell, and uses Z3 as a backend solver to fill the role of `Verify`. As previously mentioned, invariants are provided by the programmer via annotations in the code. Input to ORHLE consists of a collection of FunIMP programs, a declaration of how many copies of each program should be included in the universal and existential contexts, and a collection of function specifications expressed using the SMT-LIB2 format. Functions can have both universal and existential specifications, with the latter containing declarations of choice variables. Appendix G has example ORHLE input listings. ORHLE outputs a set of verification conditions along with a success or failure message. When a property fails to verify, ORHLE outputs a falsifying model.

Our evaluation addresses the following questions:

(R1) Is RHLE *expressive* enough to represent a variety interesting properties?
(R2) Is our approach *effective*, that is, can it be used to verify or invalidate relational assertions about a diverse corpus of programs?
(R3) Is it possible to realize an *efficient* implementation of our verification approach which return results within a reasonable time frame?

To answer these questions, we have developed a suite of 42 programs over 5 kinds of relational specifications drawn from the literature. We have also compiled an additional set of 12 benchmarks over two non-relational existentially quantified properties in order to evaluate similar questions about the non-relational existential logic from Section 4. Both sets of benchmarks contain a mix of valid and invalid properties.

Our benchmarks for the non-relational existential logic from Section 4 fall into two categories:

Winning Strategy Programs in this category play a simplified version of the card game twenty-one. Players start with two cards valued between 1 and 10, and can then request any number of additional cards. The goal is to get a hand value as close to 21 as possible without going over. The property of interest is whether an algorithmic strategy for this game permits the *possibility* of achieving the maximum hand value of 21 given any starting hand.

Branching Time Properties Our next set of benchmarks are taken from [14], which considered verification of properties of single programs expressed in CTL. The programs in this category are adaptations of the subset of those benchmarks which assert the existence of desirable final states and are thus expressible in RHLE.

Our set of relational benchmarks cover program refinement in addition to:

Noninterference Generalized noninterference is a possibilistic information security property which ensures that programs do not leak knowledge about high-security state via low-security outputs. Our formalization of this property is based on Mclean [24] and requires that, for any execution of a program p whose state is divided into high security p_H and low security p_L partitions, any other starting state with the same initial low partition can potentially yield the same final low partition, regardless of the high partition.

Delimited Release Delimited release is a relaxation of generalized noninterference which allows for limited information about secure state to be released. For example, given a confidential list of employee salaries, it may be acceptable to publicize the average salary as long as no other salary information is leaked. We formulate delimited release as a noninterference property with an additional condition requiring that the programs agree on the values of the released information. For the previous example, we would add a precondition asserting the average salary across all executions is equal.

Parameter Usage Our parameter usage benchmarks check whether a function parameter is semantically unused, in that the existence of the parameter does not affect the program's reachable final states. For example, the `flag` parameter in `f(flag) = if flag then return 1 else return 1` is syntactically used in `f`, even affecting its control flow, but does not have any effect on `f`'s possible outputs; we therefore consider `flag` to be semantically unused. For an n-ary function $f(p_1, \ldots, p_n)$, we say parameter p_i is semantically unused if

$$\langle v_i \neq w_i \wedge \bigwedge_{j \neq i} v_j = w_j \rangle \; a := f(v_1, \ldots, v_n) \sim_{\exists} b := f(w_1, \ldots, w_n) \; \langle a = b \rangle$$

Flaky Tests Tests of program behavior which can nondeterministically pass or fail pose a significant hazard as they can trigger false alarms or allow regressions to go undetected. We modeled representative nondeterministic tests in FunIMP based on examples from The Illinois Dataset of Flaky Tests (IDoFT)[33, 22], framing flakiness as a ∀∃ property containing only existential executions. We consider a test verifiably flaky when there exists both a test execution that succeeds and one that fails. We model nondeterminsitic system behavior (e.g., `getCurrentTimeMs()` or the results of network calls) as function calls. For example, to model the imprecision of thread `sleeps`, we give the verifier leeway to sleep within a ±20 ms window around the requested interval: $\mathsf{ax}_{\exists}\ \mathsf{sleep}(\text{interval}, \text{currentTime})\ [\text{sleepTime}]\ \{0 \leq \text{sleepTime} \ \wedge \ \text{interval} - 20 \leq \text{sleepTime} \leq \text{interval} + 20\}\ \{\rho = \text{currentTime} + \text{sleepTime}\}$.

The variety of properties we were able to represent in ORHLE provides evidence that it is sufficiently expressive (R1). To show that ORHLE is both

Property	**Shape**	**Pos**	**Neg**	**Unk**	**Med**(ms)	**Max**(ms)
Delimited Release	$\forall p_1 \exists p_2$	7	6	0	213	248
Flaky Tests	$\exists p_1 p_2$	2	0	0	221	230
Generalized Noninterference	$\forall p_1 \exists p_2$	5	5	1	214	232
Parameter Usage	$\forall p_1 \exists p_2$	4	2	0	210	215
Program Refinement	$\forall p_1 \exists p_2$	4	4	1	212	1349
Winning Strategy	$\exists p$	1	2	0	214	219
Branching Time	$\exists p$	7	2	0	212	219

Fig. 6. ORHLE verification results over a set of relational and non-relational properties. The **Shape** column gives the execution quantification pattern for the property; each property is of the form $\forall p_0 \dots p_n \exists q_o \dots q_n$, where p_i's and q_i's are (possibly empty) sets of executions. The **Pos** and **Neg** columns give the number of benchmarks over which the property holds or does not hold, respectively. The **Unk** column gives the number of benchmarks whose verification conditions could not be decided by the SMT solver. The **Med** and **Max** columns give (respectively) the median and maximum verification times in milliseconds over each set of benchmarks.

effective and efficient (R2)-(R3), we have used it to verify and/or invalidate examples of the benchmark properties described above. All of these experiments were done using an Intel Core i7-6700K CPU with 8 4GHz cores. Figure 6 presents the results of these experiments. ORHLE yielded the expected verification result in all cases except for one noninterference and one refinement benchmark, where the backing SMT solver (Z3) was unable to determine the validity of the verification conditions. While most benchmarks' verification conditions fell within the theory of linear integer arithmetic, verification conditions fell in a non-decidable fragment of arithmetic in both benchmarks where Z3 was unable to decide validity. One of these undecidable instances accounts for the outlier maximum verification time in the refinement benchmarks. Overall, these results offer evidence that ORHLE is both effective and efficient for verifying a variety of existential and ∀∃ properties.

7 Related Work

Relational Program Logics Relational program logics are a common approach to verifying relational specifications. Relational Hoare Logic [9] (RHL) was one of the first examples of these logics, and is capable of proving 2-safety properties. Relational Higher-order Logic [2] is a higher-order relational logic for reasoning about higher-order functional programs expressed in a simply-typed λ-calculus. Probabilistic RHL [8] is a logic for reasoning about probabilistic programs in order to prove security properties of cryptographic schemes. The relational logic closest to RHLE is Cartesian Hoare Logic [34] (CHL) developed by Sousa and Dillig. This logic which provides an axiomatic system for reasoning about k-safety hyperproperties along with an automatic verification algorithm. RHLE can be thought of as an extension of CHL for reasoning about the more general class of ∀∃ properties. Nagasamudram and Naumann [25] examine *alignment completeness* for relational Hoare logics, which classifies the ability of these logics to reason about programs in lockstep. Banerjee et al. [4] introduce a relational Hoare logic capable of reasoning about encapsulation and invariant hiding, but which is confined to 2-safety properties.

Underapproximate Program Logics Several program logics have been proposed to reason about the existence of particular executions of a single program, similar to the non-relational existential logic presented in Section 4. Reverse Hoare Logic [15] is a program logic for reasoning about reachability over single executions of programs which have access to a nondeterministic binary choice ($\sqcup$) operator. Incorrectness Logic [27] is a recent adaptation of Reverse Hoare Logic to a more realistic programming language. While these logics express the existence a satisfying start state for all satisfying end states ($\forall\sigma'\exists\sigma$), the existential logic presented in Section 4 requires there to exist a satisfying end state for all satisfying start states ($\forall\sigma\exists\sigma'$). Reverse Hoare Logic and Incorrectness Logic both reason about reachability over single executions, but properties in these logics are pure underapproximations: every state in a given postcondition must be reachable. In contrast, our reasoning over existential specifications is underapproximate *with respect to the choice variables only.* While every valid choice value must correspond to a reachable set of final states, each of these sets are still overapproximate. This feature of our existential specifications enables a natural integration with standard Hoare logics.

First-order dynamic logic [31] is a reinterpretation of Hoare logic in first-order, multi-modal logic. For a program p, the modal operators $[p]$ and $\langle p\rangle$ capture universal and existential quantification over program executions. Our universal Hoare triple $\vdash \{P\}p\{Q\}$ corresponds to $P \implies [p]Q$, and our existential Hoare triple $\vdash [P]p[Q]_\exists$ corresponds to $P \implies \langle p\rangle Q$. In contrast to RHLE, dynamic logic reasons about properties of single program executions.

Prophecy Variables Prophecy variables were originally introduced by Abadi and Lamport [1] in order to establish refinement mappings between state machines. Choice variables in our existential specifications are similar to prophecy variables in that they capture the required value of some "future" state, although we use them as part of a program logic rather than to reason about refinement mappings between state machines. Jung et al. [20] incorporate prophecy variables into a separation Hoare logic to reason about nondeterminism in concurrent programs, but differ from our approach in that the program logic operates in a non-relational setting and is designed for interactive and not automated verification.

Relational Verification The concept of a hyperproperty was originally introduced by Clarkson and Schneider [12], building on earlier work by Terauchi and Aiken [35]. The initial work discusses verification but it does not offer an algorithm; numerous program techniques have been subsequently proposed to verify hyperproperties. Product programs are an alternative approach to relational verification [5]. This approach can leverage existing non-relational verification tools and techniques when verifying the product program, but the large state space of product programs can make verification difficult in practice. Product programs have been used to verify k-safety properties and reason about noninterference and secure information flow [7, 21]. Barthe et al. [6] have developed a set of necessary conditions for “left-product programs”; these product programs can be used to verify hyperproperties outside of k-safety, including our ∀∃ properties, although the work does not address how to construct left-product programs.

Unno et al. [36] have developed a technique for verifying ∀∃ properties including program refinement, generalized noninterference, and cotermination by encoding a constraint satisfaction problem expressed using a generalization of constrained Horn clauses. The approach solves constraints using a stratified CEGIS approach, and can synthesize non-trivial alignment predicates for interleaving executions of loop bodies. This work is not based on a Hoare-style program logic, but rather develops per-property embeddings of ∀∃ verification problems in a novel adaptation of constrained Horn clauses.

There exist several modal logics which support a similar style of existential reasoning as our existential logic. Temporal logics like HyperLTL and HyperCTL [11] can be used to reason about hyperproperties, although verification tooling [10] is focused on model checking state transition systems rather than program logics. Coenen et al. [13] examine verification and synthesis of computational models using HyperLTL formulas with alternating quantifiers. Cook et al. [14] examine existential reasoning in branching-time temporal logics by way of removing state space until universal reasoning methods can be used. Lamport and Schneider [23] examine using TLA to verify ∀∃ properties including refinement and GNI. While the above approaches are capable of reasoning about the kinds of liveness properties we consider in this paper, they all focus on model checking state transition systems rather than using a Hoare-style logic to reason directly over programs as in our approach.

8 Conclusion

This paper presented RHLE, a novel relational Hoare-style program logic for reasoning about ∀∃ properties. These properties can capture a variety of interesting behaviors of multiple program executions, including program refinement and information flow properties. Key to our logic is a novel form of function specifications which constrain the set of behaviors that a valid implementation of a function *must* exhibit. We have developed an automated verification algorithm based on RHLE, and we demonstrated that an implementation of this algorithm is able to check the validity of a variety of ∀∃ properties over a benchmark suite of programs.

Bibliography

[1] M. Abadi and L. Lamport. The existence of refinement mappings. In *[1988] Proceedings. Third Annual Symposium on Logic in Computer Science*, pages 165–175, 1988.

[2] Alejandro Aguirre, Gilles Barthe, Marco Gaboardi, Deepak Garg, and Pierre-Yves Strub. A relational logic for higher-order programs. *Proc. ACM Program. Lang.*, 1(ICFP):21:1–21:29, August 2017.

[3] Andrew W. Appel. Verified software toolchain. In Gilles Barthe, editor, *Programming Languages and Systems*, pages 1–17, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[4] Anindya Banerjee, Ramana Nagasamudram, David A Naumann, and Mohammad Nikouei. A relational program logic with data abstraction and dynamic framing. *arXiv preprint arXiv:1910.14560*, 2019.

[5] Gilles Barthe, Juan Manuel Crespo, and César Kunz. Relational verification using product programs. In Michael Butler and Wolfram Schulte, editors, *FM 2011: Formal Methods*, pages 200–214, Berlin, Heidelberg, 2011. Springer Berlin Heidelberg.

[6] Gilles Barthe, Juan Manuel Crespo, and César Kunz. Beyond 2-safety: Asymmetric product programs for relational program verification. In *International Symposium on Logical Foundations of Computer Science*, pages 29–43. Springer, 2013.

[7] Gilles Barthe, Pedro R. D'Argenio, and Tamara Rezk. Secure information flow by self-composition. *Mathematical Structures in Computer Science*, 21(6):1207–1252, 2011.

[8] Gilles Barthe, Benjamin Grégoire, and Santiago Zanella Béguelin. Formal certification of code-based cryptographic proofs. *SIGPLAN Not.*, 44(1):90–101, January 2009.

[9] Nick Benton. Simple relational correctness proofs for static analyses and program transformations. In *Proceedings of the 31st ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, POPL '04, pages 14–25, New York, NY, USA, 2004. ACM.

[10] E. Clarke, O. Grumberg, and D. Long. Verification tools for finite-state concurrent systems. In J. W. de Bakker, W. P. de Roever, and G. Rozenberg, editors, *A Decade of Concurrency Reflections and Perspectives*, pages 124–175, Berlin, Heidelberg, 1994. Springer Berlin Heidelberg.

[11] Michael R Clarkson, Bernd Finkbeiner, Masoud Koleini, Kristopher K Micinski, Markus N Rabe, and César Sánchez. Temporal logics for hyperproperties. In *International Conference on Principles of Security and Trust*, pages 265–284. Springer, 2014.

[12] Michael R. Clarkson and Fred B. Schneider. Hyperproperties. *J. Comput. Secur.*, 18(6):1157–1210, September 2010.

[13] Norine Coenen, Bernd Finkbeiner, César Sánchez, and Leander Tentrup. Verifying hyperliveness. pages 121–139, 07 2019.

[14] Byron Cook and Eric Koskinen. Reasoning about nondeterminism in programs. In *Proceedings of the 34th ACM SIGPLAN conference on Programming language design and implementation*, pages 219–230, 2013.

[15] Edsko de Vries and Vasileios Koutavas. Reverse hoare logic. In *Proceedings of the 9th International Conference on Software Engineering and Formal Methods*, SEFM'11, page 155–171, Berlin, Heidelberg, 2011. Springer-Verlag.

[16] Isil Dillig, Thomas Dillig, Boyang Li, and Ken McMillan. Inductive invariant generation via abductive inference. In *Proceedings of the 2013 ACM SIGPLAN International Conference on Object Oriented Programming Systems Languages and Applications*, OOPSLA '13, page 443–456, New York, NY, USA, 2013. Association for Computing Machinery.

[17] Cormac Flanagan and K. Rustan M. Leino. Houdini, an Annotation Assistant for ESC/Java. In *Proceedings of the International Symposium of Formal Methods Europe on Formal Methods for Increasing Software Productivity*, FME '01, page 500–517, Berlin, Heidelberg, 2001. Springer-Verlag.

[18] C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, October 1969.

[19] Ralf Jung, Jacques-Henri Jourdan, Robbert Krebbers, and Derek Dreyer. Rustbelt: Securing the foundations of the rust programming language. *Proc. ACM Program. Lang.*, 2(POPL), dec 2017.

[20] Ralf Jung, Rodolphe Lepigre, Gaurav Parthasarathy, Marianna Rapoport, Amin Timany, Derek Dreyer, and Bart Jacobs. The future is ours: Prophecy variables in separation logic. *Proc. ACM Program. Lang.*, 4(POPL), December 2019.

[21] Máté Kovács, Helmut Seidl, and Bernd Finkbeiner. Relational abstract interpretation for the verification of 2-hypersafety properties. pages 211–222, 11 2013.

[22] W. Lam, R. Oei, A. Shi, D. Marinov, and T. Xie. idflakies: A framework for detecting and partially classifying flaky tests. In *2019 12th IEEE Conference on Software Testing, Validation and Verification (ICST)*, pages 312–322, 2019.

[23] Leslie Lamport and Fred B. Schneider. Verifying Hyperproperties With TLA. In *2021 IEEE 34th Computer Security Foundations Symposium (CSF)*, pages 1–16, June 2021. ISSN: 2374-8303.

[24] John McLean. A general theory of composition for a class of "possibilistic" properties. *IEEE Trans. Softw. Eng.*, 22(1):53–67, January 1996.

[25] Ramana Nagasamudram and David A. Naumann. Alignment completeness for relational hoare logics. In *2021 36th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, pages 1–13, 2021.

[26] Peter W. O'Hearn. Resources, concurrency, and local reasoning. *Theoretical Computer Science*, 375(1):271–307, 2007. Festschrift for John C. Reynolds's 70th birthday.

[27] Peter W. O'Hearn. Incorrectness logic. *Proc. ACM Program. Lang.*, 4(POPL), December 2019.

[28] Saswat Padhi, Rahul Sharma, and Todd Millstein. Data-driven precondition inference with learned features. *ACM SIGPLAN Notices*, 51(6):42–56, 2016.

[29] Saswat Padhi, Rahul Sharma, and Todd Millstein. LoopInvGen: A Loop Invariant Generator based on Precondition Inference, 2017.

[30] Peter Poetzsch-Heffter, Arndand Müller. A Programming Logic for Sequential Java. In S. Doaitse Swierstra, editor, *Programming Languages and Systems*, pages 162–176, Berlin, Heidelberg, 1999. Springer Berlin Heidelberg.

[31] Vaughan R Pratt. Semantical consideration on Floyd-Hoare logic. In *17th Annual Symposium on Foundations of Computer Science (sfcs 1976)*, pages 109–121. IEEE, 1976.

[32] J.C. Reynolds. Separation logic: a logic for shared mutable data structures. In *Proceedings 17th Annual IEEE Symposium on Logic in Computer Science*, pages 55–74, 2002.

[33] A. Shi, A. Gyori, O. Legunsen, and D. Marinov. Detecting assumptions on deterministic implementations of non-deterministic specifications. In *2016 IEEE International Conference on Software Testing, Verification and Validation (ICST)*, pages 80–90, 2016.

[34] Marcelo Sousa and Isil Dillig. Cartesian hoare logic for verifying k-safety properties. In *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation*, PLDI '16, pages 57–69, New York, NY, USA, 2016. ACM.

[35] Tachio Terauchi and Alex Aiken. Secure information flow as a safety problem. In Chris Hankin and Igor Siveroni, editors, *Static Analysis*, pages 352–367, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.

[36] Hiroshi Unno, Tachio Terauchi, and Eric Koskinen. Constraint-Based Relational Verification. In Alexandra Silva and K. Rustan M. Leino, editors, *Computer Aided Verification*, Lecture Notes in Computer Science, pages 742–766, Cham, 2021. Springer International Publishing.

A Semantics of FunIMP

The semantics of FunIMP is given as a big-step reduction relation from initial to final states. This relation is parameterized over an *implementation* context $I \in \mathcal{N} \to FD$, a partial mapping from function names to definitions. FunIMP program states, $\sigma \in \mathcal{V} \to \mathbb{N}$, are mappings from variables to their current value. The reduction relation is also parameterized over an interpretation used to determine the validity of assertions; we write $\sigma \models P$ to denote that the assertion P holds in state σ. We condense sequences of repeated expressions in similar way to function arguments, writing the sequence $I \vdash \sigma, a_1 \Downarrow v_1 \;\cdots\; I \vdash \sigma, a_n \Downarrow v_n$ as $I \vdash \sigma, \overline{a} \Downarrow \overline{v}$ and $[x_1 \mapsto v_1, \ldots, x_n \mapsto v_n]$ as $[\overline{x} \mapsto \overline{v}]$, for example.

The evaluation rules of FunIMP are presented in Figure 7.

$$\frac{}{I \vdash \sigma, \mathsf{skip} \Downarrow \sigma}\ \textsc{ESkip} \qquad \frac{\sigma, a \Downarrow v}{I \vdash \sigma, x := a \Downarrow [x \mapsto v]\sigma}\ \textsc{EAssgn}$$

$$\frac{}{I \vdash \sigma, x ::= \mathsf{havoc} \Downarrow [x \mapsto v]\sigma}\ \textsc{EHavoc} \qquad \frac{I \vdash \sigma, c_1 \Downarrow \sigma' \qquad I \vdash \sigma', c_1 \Downarrow \sigma''}{I \vdash \sigma, c_1;\ c_2 \Downarrow \sigma''}\ \textsc{ESeq}$$

$$\frac{\sigma, b \Downarrow \mathsf{true} \qquad I \vdash \sigma, c_1 \Downarrow \sigma'}{I \vdash \sigma, \mathsf{if}\ b\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2 \Downarrow \sigma'}\ \textsc{ECondT} \qquad \frac{\sigma, b \Downarrow \bot \qquad I \vdash \sigma, c_2 \Downarrow \sigma'}{I \vdash \sigma, \mathsf{if}\ b\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2 \Downarrow \sigma'}\ \textsc{ECondF}$$

$$\frac{\begin{array}{c}\sigma, b \Downarrow \mathsf{true} \qquad I \vdash \sigma, c \Downarrow \sigma' \\ I \vdash \sigma', \mathsf{while}\ b\ \mathsf{do}\ c\ \mathsf{end} \Downarrow \sigma''\end{array}}{I \vdash \sigma, \mathsf{while}\ b\ \mathsf{do}\ c\ \mathsf{end} \Downarrow \sigma''}\ \textsc{EWhileT} \qquad \frac{\sigma, b \Downarrow \mathsf{false}}{I \vdash \sigma, \mathsf{while}\ b\ \mathsf{do}\ c\ \mathsf{end} \Downarrow \sigma}\ \textsc{EWhileF}$$

$$\frac{\begin{array}{c}I(f) = \textbf{def}\ f(\overline{x})\,\{s; \textbf{return}\ e\} \\ I \vdash \sigma, \overline{a} \Downarrow \overline{v} \qquad I \vdash [\overline{x} \mapsto \overline{v}], s \Downarrow \sigma' \qquad I \vdash \sigma', e \Downarrow r\end{array}}{I \vdash \sigma, y := f(\overline{a}) \Downarrow [y \mapsto r]\sigma}\ \textsc{ECall}$$

Fig. 7. Big-step evaluation relation of FunIMP with a concrete implementation context.

A.1 Overapproximate Executions Semantics

The big-step operational semantics for overapproximate evaluation is given in Figure 8. These semantics are nearly identical to the evaluation semantics over concrete implementation contexts given in Figure 7, but is instead parameterized over a universal specification context $S_\forall \in \mathcal{N} \to FA$ and replaces the ECall rule with the $\textsc{ECall}_\forall$ rule. The latter rule allows a call to a universally specified function to step to any state with a return value consistent with the function's specification.

$$\frac{}{S_\forall \vdash \sigma, \mathsf{skip} \Downarrow_\forall \sigma}\ \textsc{ESkip}_\forall \qquad\qquad \frac{\sigma, a \Downarrow_\forall v}{S_\forall \vdash \sigma, x := a \Downarrow_\forall [x \mapsto v]\sigma}\ \textsc{EAssgn}_\forall$$

$$\frac{}{S_\forall \vdash \sigma, x := \mathsf{havoc} \Downarrow_\forall [x \mapsto v]\sigma}\ \textsc{EHavoc}_\forall$$

$$\frac{S_\forall \vdash \sigma, c_1 \Downarrow_\forall \sigma' \qquad S_\forall \vdash \sigma', c_1 \Downarrow_\forall \sigma''}{S_\forall \vdash \sigma, c_1;\ c_2 \Downarrow_\forall \sigma''}\ \textsc{ESeq}_\forall$$

$$\frac{\sigma, b \Downarrow_\forall \mathsf{true} \qquad S_\forall \vdash \sigma, c_1 \Downarrow_\forall \sigma'}{S_\forall \vdash \sigma, \mathsf{if}\ b\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2 \Downarrow_\forall \sigma'}\ \textsc{ECondT}_\forall$$

$$\frac{\sigma, b \Downarrow_\forall \bot \qquad S_\forall \vdash \sigma, c_2 \Downarrow_\forall \sigma'}{S_\forall \vdash \sigma, \mathsf{if}\ b\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2 \Downarrow_\forall \sigma'}\ \textsc{ECondF}_\forall$$

$$\frac{\begin{array}{c}\sigma, b \Downarrow_\forall \mathsf{true} \qquad S_\forall \vdash \sigma, c \Downarrow_\forall \sigma' \\ S_\forall \vdash \sigma', \mathsf{while}\ b\ \mathsf{do}\ c\ \mathsf{end} \Downarrow_\forall \sigma''\end{array}}{S_\forall \vdash \sigma, \mathsf{while}\ b\ \mathsf{do}\ c\ \mathsf{end} \Downarrow_\forall \sigma''}\ \textsc{EWhileT}_\forall$$

$$\frac{\sigma, b \Downarrow_\forall \mathsf{false}}{S_\forall \vdash \sigma, \mathsf{while}\ b\ \mathsf{do}\ c\ \mathsf{end} \Downarrow_\forall \sigma}\ \textsc{EWhileF}_\forall$$

$$\frac{\begin{array}{c}S_\forall(f) = \mathsf{ax}_\forall\ f(\overline{x})\,\{P\}\,\{Q\} \\ S_\forall \vdash \sigma, \overline{a} \Downarrow_\forall \overline{v} \qquad [\overline{x} \mapsto \overline{v}] \models P \qquad [\rho \mapsto r, \overline{x} \mapsto \overline{v}] \models Q\end{array}}{S_\forall \vdash \sigma, y := f(\overline{a}) \Downarrow_\forall [y \mapsto r]\sigma}\ \textsc{ECall}_{\forall 1}$$

$$\frac{S_\forall(f) = \mathsf{ax}_\forall\ f(\overline{x})\,\{P\}\,\{Q\} \qquad S_\forall \vdash \sigma, \overline{a} \Downarrow_\forall \overline{v} \qquad [\overline{x} \mapsto \overline{v}] \not\models P}{S_\forall \vdash \sigma, y := f(\overline{a}) \Downarrow_\forall [y \mapsto r]\sigma}\ \textsc{ECall}_{\forall 2}$$

Fig. 8. Overapproximate execution semantics of FunIMP with a universal specification context.

B Universal Hoare Logic

$$\frac{\models P \implies P' \qquad \models Q' \implies Q \qquad S_\forall \vdash \{P'\}\ c\ \{Q'\}}{S_\forall \vdash \{P\}\ c\ \{Q\}}\ \forall\textsc{Conseq}$$

$$\frac{}{S_\forall \vdash \{P\}\ \mathsf{skip}\ \{P\}}\ \forall\textsc{Skip} \qquad \frac{}{S_\forall \vdash \{P[a/x]\}\ x := a\ \{P\}}\ \forall\textsc{Assgn}$$

$$\frac{}{S_\forall \vdash \{\forall v.P[v/x]\}\ x ::= \mathsf{havoc}\ \{P\}}\ \forall\textsc{Havoc}$$

$$\frac{S_\forall \vdash \{P\}\ c_1\ \{P'\} \qquad S_\forall \vdash \{P'\}\ c_2\ \{Q\}}{S_\forall \vdash \{P\}\ c_1;\ c_2\ \{Q\}}\ \forall\textsc{Seq}$$

$$\frac{S_\forall \vdash \{P \land b\}\ c_1\ \{Q\} \qquad S_\forall \vdash \{P \land \neg b\}\ c_2\ \{Q\}}{S_\forall \vdash \{P\}\ \mathsf{if}\ b\ \mathsf{then}\ c_1\ \mathsf{else}\ c_2\ \{Q\}}\ \forall\textsc{Cond}$$

$$\frac{S_\forall \vdash \{P\ \land\ b\}\ c\ \{P\}}{S_\forall \vdash \{P\}\ \mathsf{while}\ b\ \mathsf{do}\ c\ \mathsf{end}\ \{P\ \land\ \neg b\}}\ \forall\textsc{While}$$

$$\frac{S(f) = \mathsf{ax}\ f(\overline{x})\{P\}\{Q\}}{S_\forall \vdash \left\{ P[\overline{a}/\overline{x}] \land \quad \forall v.Q[v/\rho; \overline{a}/\overline{x}] \quad \implies Q[v/y] \right\}\ y := f(\overline{a})\ \{Q\}}\ \forall\textsc{Spec}$$

Fig. 9. Proof rules for a universal Hoare logic for FunIMP.

C Existential Hoare Logic

$$\frac{\models P \implies P' \qquad \models Q' \implies Q \qquad S_\exists \vdash [P']\ s\ [Q']_\exists}{S_\exists \vdash [P]\ s\ [Q]_\exists}\ \exists\textsc{Conseq}$$

$$\frac{}{S_\exists \vdash [P]\ \mathsf{skip}\ [P]_\exists}\ \exists\textsc{Skip} \qquad\qquad \frac{}{S_\exists \vdash [Q[a/x]]\ x := a\ [Q]_\exists}\ \exists\textsc{Assgn}$$

$$\frac{}{S_\exists \vdash [\exists v.\ Q[v/x]]\ x := \mathsf{havoc}\ [Q]_\exists}\ \exists\textsc{Havoc}$$

$$\frac{S_\exists \vdash [P]\ s_1\ [P']_\exists \qquad S_\exists \vdash [P']\ s_2\ [Q]_\exists}{S_\exists \vdash [P]\ s_1;\ s_2\ [Q]_\exists}\ \exists\textsc{Seq}$$

$$\frac{S_\exists \vdash [P \land b]\ s_1\ [Q]_\exists \qquad S_\exists \vdash [P \land \neg b]\ s_2\ [Q]_\exists}{S_\exists \vdash [P]\ \mathsf{if}\ b\ \mathsf{then}\ s_1\ \mathsf{else}\ s_2\ [Q]_\exists}\ \exists\textsc{Cond}$$

$$\frac{R\ \textsf{is well-founded} \qquad S_\exists \vdash [P\ \land\ b\ \land\ M\,a]\ s\ [P\ \land\ \exists a'.\, M\,a'\ \land\ a'\,R\,a]_\exists}{S_\exists \vdash [P \land \exists a.\, M\,a]\ \mathsf{while}\ b\ \mathsf{do}\ s\ \mathsf{end}\ [P\ \land\ \neg b]_\exists}\ \exists\textsc{While}$$

$$\frac{S_\exists(f) = \mathsf{ax}_\exists f(\overline{x})\ [\overline{c}]\ \{P\}\ \{Q\}}{S_\exists \vdash \begin{bmatrix} \exists\overline{k}.\ ([\overline{x} \mapsto \overline{a}] \models P[\overline{k}/\overline{c}] \\ \land\ \ \exists v.[\rho \mapsto v, \overline{x} \mapsto \overline{a}] \models Q[\overline{k}/\overline{c}] \\ \land\ \ \forall v.[\rho \mapsto v, \overline{x} \mapsto \overline{a}] \models Q[\overline{k}/\overline{c}] \\ \implies R[v/y]) \end{bmatrix}\ y := f(\overline{a})\ [R]_\exists}\ \exists\textsc{Spec}$$

Fig. 10. Existential Hoare logic rules.

D Synchronous Rules

$$\dfrac{S_\forall, S_\exists \vdash \langle \Phi \rangle\ \overline{\mathsf{skip};\ p_\forall} \sim_\exists \overline{\mathsf{skip};\ p_\exists}\ \langle \Psi \rangle}{S_\forall, S_\exists \vdash \langle \Phi \rangle\ \overline{p_\forall} \sim_\exists \overline{p_\exists}\ \langle \Psi \rangle}\ \textsc{SkipIntroL}$$

$$\dfrac{S_\forall, S_\exists \vdash \langle \Phi \rangle\ \overline{s_\forall} \sim_\exists \overline{s_\exists}\ \langle \chi \rangle \qquad S_\forall, S_\exists \vdash \langle \chi \rangle\ \overline{s'_\forall} \sim_\exists \overline{s'_\exists}\ \langle \Psi \rangle}{S_\forall, S_\exists \vdash \langle \Phi \rangle\ \overline{s_\forall;\ s'_\forall} \sim_\exists \overline{s_\exists;\ s'_\exists}\ \langle \Psi \rangle}\ \textsc{SyncSeq}$$

$$\dfrac{\begin{array}{c} S_\forall, S_\exists \vdash \langle \mathbb{I} \wedge \bigwedge\limits_{0 \le i \le n} b_i \rangle\ s_0 \circledast \cdots \circledast s_k \sim_\exists s_{k+1} \circledast \cdots \circledast s_n\ \langle \mathbb{I} \rangle \\ \mathbb{I} \wedge \bigwedge\limits_{0 \le i \le n} \neg b_i \implies \Psi \qquad \mathbb{I} \wedge \neg \bigwedge\limits_{0 \le i \le n} b_i \implies \bigwedge\limits_{0 \le i \le n} \neg b_i \end{array}}{\begin{array}{c} S_\forall, S_\exists \vdash \langle \mathbb{I} \rangle\ \mathsf{while}\ b_0\ \mathsf{do}\ s_0\ \mathsf{end} \circledast \cdots \circledast \mathsf{while}\ b_k\ \mathsf{do}\ s_k\ \mathsf{end} \\ \sim_\exists \mathsf{while}\ b_{k+1}\ \mathsf{do}\ s_{k+1}\ \mathsf{end} \circledast \cdots \circledast \mathsf{while}\ b_n\ \mathsf{do}\ s_n\ \mathsf{end}\ \langle \Psi \rangle \end{array}}\ \textsc{SyncLoops}$$

$$\dfrac{\begin{array}{c} R\ \text{is well-founded} \\ S_\forall, S_\exists \vdash \langle \mathbb{I} \wedge \bigwedge\limits_{0 \le i \le n} b_i \wedge M\ a \rangle\ \overline{\mathsf{skip}} \sim_\exists s_0 \circledast \cdots \circledast s_n\ \langle \mathbb{I} \wedge \exists a'.\ M\ a' \ \wedge\ a'\ R\ a \rangle \\ \mathbb{I} \wedge \bigwedge\limits_{0 \le i \le n} \neg b_i \implies \Psi \qquad \mathbb{I} \wedge \neg \bigwedge\limits_{0 \le i \le n} b_i \implies \bigwedge\limits_{0 \le i \le n} \neg b_i \end{array}}{S_\forall, S_\exists \vdash \langle \mathbb{I} \wedge \exists a.\ M\ a \rangle\ \overline{\mathsf{skip}} \qquad \sim_\exists \mathsf{while}\ b_0\ \mathsf{do}\ s_0\ \mathsf{end} \circledast \cdots \circledast \mathsf{while}\ b_n\ \mathsf{do}\ s_n\ \mathsf{end}\ \langle \Psi \rangle}\ \textsc{SyncLoops}_\exists$$

$$\dfrac{\begin{array}{c} S_\forall, S_\exists \vdash \langle \mathbb{I} \wedge \bigwedge\limits_{0 \le i \le n} b_i \rangle\ s_0 \circledast \cdots \circledast s_n \sim_\exists \overline{\mathsf{skip}}\ \langle \mathbb{I} \rangle \\ \mathbb{I} \wedge \bigwedge\limits_{0 \le i \le n} \neg b_i \implies \Psi \qquad \mathbb{I} \wedge \neg \bigwedge\limits_{0 \le i \le n} b_i \implies \bigwedge\limits_{0 \le i \le n} \neg b_i \end{array}}{S_\forall, S_\exists \vdash \langle \mathbb{I} \wedge \exists a.\ M\ a \rangle\ \mathsf{while}\ b_0\ \mathsf{do}\ s_0\ \mathsf{end} \circledast \cdots \circledast \mathsf{while}\ b_n\ \mathsf{do}\ s_n\ \mathsf{end} \sim_\exists \overline{\mathsf{skip}}\ \langle \Psi \rangle}\ \textsc{SyncLoops}_\forall$$

Fig. 11. Synchronous RHLE proof rules.

E Proofs

Theorem 4. *When run under an implementation context I that is $\forall$-compatible with specification context $S_\forall$ with an initial state σ, a program p will either diverge or evaluate to a state σ' which is also the result of one of its overapproximate executions under $S_\forall$:*

$$I \models_\forall S_\forall \quad \wedge \quad I \vdash \sigma, p \Downarrow \sigma' \quad \implies \quad S_\forall \vdash \sigma, p \Downarrow_\forall \sigma'$$

Proof. By induction over the derivation of $I \vdash \sigma, p \Downarrow \sigma'$. The only interesting case is ECall, where we must consider whether σ meets the precondition of f

in $S_\forall$. If not, the proof is immediate from ECall$_{\forall 2}$. If so, the proof follows from the fact that the definition of f in I is compatible with its specification in $S_\forall$ and ECall$_{\forall 1}$.

Theorem 5. *If there is an underapproximate evaluation of program p to a set of states Σ from an initial state σ under $S_\exists$, then p must terminate in at least one final state $\sigma' \in \Sigma$ from σ under every implementation context I that is $\exists$-compatible with $S_\exists$:*

$$S_\exists \vdash \sigma, p \Downarrow_\exists \Sigma \quad \wedge \quad I \models_\exists S_\exists \quad \implies \quad \exists \sigma'. I \vdash \sigma, p \Downarrow \sigma' \; \wedge \; \sigma' \in \Sigma$$

Proof. By induction over the derivation of $S_\exists \vdash \sigma, p \Downarrow_\exists \Sigma$. Once again, the interesting case is ECall$_\exists$, which follows immediate from the fact that I is $\exists$-compatible with $S_\exists$.

Theorem 6 (RHLE is Sound). *Suppose $S_\forall, S_\exists \vdash \langle \Phi \rangle \; \overline{p_\forall} \sim_\exists \overline{p_\exists} \; \langle \Psi \rangle$. Then, for any function context I compatible with $S_\forall$ and $S_\exists$, any set of initial states $\overline{\sigma_\forall}$ and $\overline{\sigma_\exists}$ satisfying Φ, and every collection of final states $\overline{\sigma'_\forall}$ of $\overline{p_\forall}$, there must exist a collection of final states produced by $\overline{p_\exists}$ that, together with $\overline{\sigma'_\forall}$ satisfies the relational post-condition Ψ:*

$$\begin{aligned}
&\; S_\forall, S_\exists \vdash \langle \Phi \rangle \; \overline{p_\forall} \sim_\exists \overline{p_\exists} \; \langle \Psi \rangle \\
&\wedge \forall I. \, I \models S_\forall \wedge I \models S_\exists \\
&\wedge \forall \overline{\sigma_\forall} \, \overline{\sigma_\exists}. \, \overline{\sigma_\forall}, \overline{\sigma_\exists} \models \Phi \\
&\wedge \forall \overline{\sigma'_\forall}. \, I \vdash \overline{\sigma_\forall}, \overline{p_\forall} \; \Downarrow \; \overline{\sigma'_\forall} \implies \\
&\; \exists \overline{\sigma'_\exists}. \, I \vdash \sigma_\exists, \overline{p_\exists} \; \Downarrow \; \overline{\sigma'_\exists} \; \wedge \; \overline{\sigma'_\forall}, \; \overline{\sigma'_\exists} \models \Psi
\end{aligned}$$

Proof. We first prove a stronger property by induction on the triple $S_\forall, S_\exists \vdash \langle \Phi \rangle \; \overline{p_\forall} \sim_\exists \overline{p_\exists} \; \langle \Psi \rangle$: namely, that there exist appropriate existential executions of $p_\exists$ for every collection of final states of $p_\forall$ produced by an overapproximate execution, for any set of initial states satisfying the precondition Φ:

$$\begin{aligned}
&\forall \overline{\sigma_\forall} \, \overline{\sigma_\exists}. \, \overline{\sigma_\forall} \, \overline{\sigma_\exists} \models \Phi \quad \wedge \quad \forall \overline{\sigma'_\forall}. \, S_\forall \vdash \overline{\sigma_\forall}, \overline{p_\forall} \; \Downarrow_\forall \; \overline{\sigma'_\forall} \implies \\
&S_\exists \vdash \sigma_\exists, \overline{p_\exists} \Downarrow_\exists \{ \overline{\sigma'_\exists} \mid \overline{\sigma'_\forall}, \overline{\sigma'_\exists} \models \Psi \}
\end{aligned} \tag{1}$$

By Theorem 1, the fact that I is $\forall$-compatible with $S_\forall$, and our assumption that $I \vdash \overline{\sigma_\forall}, \overline{p_\forall} \Downarrow \overline{\sigma'_\forall}$, it follows that:

$$S_\forall \vdash \overline{\sigma_\forall}, \overline{p_\forall} \Downarrow_\forall \overline{\sigma'_\forall} \tag{2}$$

Armed with (1) and (2) and the assumption that I is $\exists$-compatible with $S_\exists$, by Theorem 2 we can conclude the desired result, i.e. $\exists \overline{\sigma'_\exists}. I \vdash \sigma_\exists, \overline{p_\exists} \Downarrow \overline{\sigma'_\exists} \wedge \overline{\sigma'_\forall}, \; \overline{\sigma'_\exists} \models \Psi$.

[RD: TODO: Add statements of soundness of the non-relational universal and existential Hoare logics to the appendix.]

F Verification Example

To illustrate the operation of `VCGen`, consider proving the following simple refinement assertion, where the contexts $S_\forall$ and $S_\exists$ contain the specifications for randB from Example 1:

$$S_\forall, S_\exists \vdash \langle \top \rangle\ y_1 := \mathsf{randB}(4) \sim_\exists y_2 := \mathsf{randB}(10)\ \langle y_1 = y_2 \rangle$$

`RHLEVerify` begins by calling `VCGen` with:

$$p_\forall \equiv \{\mathsf{skip}; y_1 := \mathsf{randB}(4)\} \qquad p_\exists \equiv \{\mathsf{skip}; y_2 := \mathsf{randB}(10)\} \qquad \overline{\Psi} \equiv (\varnothing, \varnothing, y_1 = y_2)$$

`VCGen` matches the randB call in $p_\forall$ (line 5), and recurses with the new postcondition built by $wp_\forall$ (omitting the trivial precondition for brevity):

$$p_\forall \equiv \{\mathsf{skip}\} \qquad p_\exists \equiv \{y_2 := \mathsf{randB}(10)\} \qquad \overline{\Psi} : (\{v_1\}, \varnothing, 0 \le v_1 < 4 \implies v_1 = y_2)$$

`VCGen` now chooses the existential call to randB (line 7), and recurses again with a postcondition built by $wp_\exists$:

$$p_\forall \equiv \{\mathsf{skip}\} \qquad p_\exists \equiv \{\mathsf{skip}\} \qquad \overline{\Psi} : (\{v_1\}, \{v_2\}, 0 \le v_2 < 10 \wedge (0 \le v_1 < 4 \implies v_1 = v_2))$$

Since both programs are now skip, `VCGen` terminates, returning $(\{v_1\}, \{v_2\}, 0 \le v_2 < 10 \wedge (0 \le v_1 < 4 \implies v_1 = v_2))$. `RHLEVerify` uses this formula to construct the following query:

$$\forall v_1\ \exists v_2.\ 0 \le v_2 < 10 \wedge (0 \le v_1 < 4 \implies v_1 = v_2)$$

which it hands off to the `Verify`. Note how this formula encodes the essence of the $\forall\exists$ question posed by the original triple; for all allowed randB return values v_1 in the universal execution, we want to know if there exists an allowed instantiation of the choice variable v_2 in the existential execution that brings both programs to the same final state. In this case, the $\forall\exists$ formula is valid, and verification succeeds.

G Example ORHLE Input

The following listing verifies a noninterference property, namely that the program never leaks any information about the variable high. Note that the underapproximation of `flipCoin` is required. If linked to a `flipCoin` implementation that always returns 0, for example, attackers could always know whether or not the initial value of `low` was less than `high` by observing low.

```
forall: run[1];
exists: run[2];

pre:  (= run!1!low run!2!low);
post: (= run!1!low run!2!low);
```

```
aspecs:
  flipCoin() {
    pre: true;
    post: (or (= ret! 0) (= ret! 1));
  }

especs:
  flipCoin() {
    choiceVars: n;
    pre: (or (= n 0) (= n 1));
    post: (= ret! n);
  }

fun run(high, low) {
  if (low < high) then
    low := 0;
  else
    low := 1;
  end
  flip := call flipCoin();
  if (flip == 0) then
    low := 1 − low;
  endif
}
```

Conversely, ORHLE identifies a violation of noninterference in the listing below. The program might leak the value of high, depending on the outcome of `flipCoin`.

```
forall: run[1];
exists: run[2];

pre: (= run!1!low run!2!low);
post: (= run!1!low run!2!low);

aspecs:
  flipCoin() {
    pre: true;
    post: (or (= ret! 0) (= ret! 1));
  }

especs:
  flipCoin() {
    choiceVars: n;
    pre: (or (= n 0) (= n 1));
    post: (= ret! n);
  }

fun run(high, low) {
  flip := call flipCoin();
  if (flip == 0) then
    low := high + low;
  endif
}
```