

Validating Library Usage Interactively^{*}

William R. Harris, Guoliang Jin, Shan Lu, and Somesh Jha

University of Wisconsin, Madison, WI, USA
{ wrharris, aliang, shanlu, jha }@cs.wisc.edu

Abstract. Programmers who develop large, mature applications often want to optimize the performance of their program without changing its semantics. They often do so by changing how their program invokes a library function or a function implemented in another module of the program. Unfortunately, once a programmer makes such an optimization, it is difficult for him to validate that the optimization does not change the semantics of the original program, because the original and optimized programs are equivalent only due to subtle, implicit assumptions about library functions called by the programs.

In this work, we present an interactive program analysis that a programmer can apply to validate that his optimization does not change his program's semantics. Our analysis casts the problem of validating an optimization as an abductive inference problem in the context of checking program equivalence. Our analysis solves the abductive equivalence problem by interacting with the programmer so that the programmer implements a solver for a logical theory that models library functions invoked by the program. We have used our analysis to validate optimizations of real-world, mature applications: the Apache software suite, the Mozilla Suite, and the MySQL database.

Keywords: abductive reasoning, program equivalence

1 Introduction

Application developers often modify a program to produce a new program that executes faster than, but is semantically equivalent to, the original program. After a developer modifies his program, he can determine with high confidence whether the modified program executes faster than the original program by measuring the performance of the original and modified program on a set of performance benchmarks. Unfortunately, it is significantly harder for the developer to determine that the modified program is semantically equivalent to the original program.

^{*} Supported, in part, by DARPA and AFRL under contract FA8650-10-C-7088. The views, opinions, and/or findings contained herein are those of the authors and should not be interpreted as representing the official views or policies, either expressed or implied, of DARPA or the Department of Defense. Also supported by NSF under grants CCF-1054616, CCF-1217582, and by a Clare Boothe Luce faculty fellowship.

Much previous work in developing correct compiler optimizations has focused on developing fully-automatic analyses that determine if two programs are equivalent [23, 25]. Unfortunately, such analyses usually require that the two programs call the same procedures with the same argument values. However, many practical optimizations modify a program to call a library function on different values, or call a different library function entirely. Such analyses cannot prove that such an optimization preserves the semantics of a program. Other analyses attempt to determine if two programs are equivalent by analyzing the programs inter-procedurally [10, 17]. Unfortunately, many practical optimizations modify calls to complex, heavily optimized library functions. Such functions may be difficult to analyze, or their source code may be unavailable.

In this work, we propose a new interactive analysis for determining under what conditions two programs are equivalent. Unlike previous work, the analysis that we propose is not fully automatic. Instead, the analysis takes as input from a programmer an original program and an optimized program, and suggests a candidate specification of the functions defined in libraries and other program modules (in this paper, we refer to all such functions as “library” functions for simplicity) called by the program that implies that the original and optimized programs are equivalent. The programmer either validates or refutes the candidate specification, and the analysis uses this validation or refutation to iteratively suggest a new sufficient specification, until the analysis finds a sufficient specification that is validated by the programmer. If a programmer accepts an invalid specification of library functions, then the analysis may incorrectly determine that the programs are equivalent. However, even if a programmer accepts an invalid specification, the analysis still generates an explicit representation of the key assumptions made by the programmer to justify the optimization. The assumptions can potentially be validated by other programmers or known techniques for verifying safety properties of programs [4, 6, 13].

There are two key challenges to developing an interactive equivalence checker. The first key challenge is to develop a checker that can construct candidate specifications about functions whose implementations may not be available, or that manipulate complex abstract datatypes, such as strings, that are difficult to reason about symbolically. The equivalence checker must find specifications describing such functions that it can soundly determine to be consistent and sufficient to prove that the original program is equivalent to the optimized program.

The second key challenge is to develop an interactive checker that queries its user with simple, non-redundant candidate specifications about the library functions that a program calls. To prove equivalence between an original and optimized program, the interactive checker must work with the user to construct a simulation relation from the state space of the original program to the state space of the optimized program. However, the checker should largely hide from the user the complexity of constructing such a simulation relation, so that the user must only ever make simple Boolean decisions determining the validity of a candidate specification of library functions.

```

string srch_strm(string s) {
L0: string sb := "";
L1: while(!find(sb, s)) {
    char c := get();
    sb := append(sb, c);
}
L2: return sb;
}

string srch_strm'(string s) {
L0': str sb := "";
    int pos := - len(s);
L1': while(pos < 0
    || !find(sub(sb, pos), s)) {
    char c := get();
    sb := append(sb, c);
    pos := len(sb) - len(s);
}
L2': return sb;
}

```

Fig. 1. An example original program `srch_strm` and its optimization `srch_strm'`, simplifications of the original and optimized programs submitted in Apache Bug #34464.

Our key insight to address the above challenges is to design the checker so that it treats the user as a solver for a theory that describes the library functions. The equivalence checker reduces the problem of checking equivalence to proving the validity of a set of formulas, using known techniques for checking equivalence [23]. To prove validity of the required formulas, the equivalence checker applies an *abductive theorem prover*, which generates an assumption over the library functions, restricted to logical combinations of equalities, that are sufficient for each formula to be valid. To generate such an assumption, the theorem prover uses an *optimistic solver* for the theory of program libraries. If the optimistic solver finds consistent assumptions sufficient to prove validity, then the equivalence checker presents the assumptions for the user to validate or refute. In other words, the optimistic solver interacts with the user to implement a “guess-and-check” solver for the theory of the libraries.

The rest of this paper is organized as follows: §2 illustrates our equivalence problem and analysis on a function and its optimization submitted in a bug report for the Apache Ant build tool. §3 presents our abductive equivalence problem and analysis in detail. §4 presents an experimental evaluation of our analysis on a set of benchmarks taken from bug reports to fix performance issues in applications. §5 discusses related work.

2 Overview

In this section, we motivate the abductive equivalence problem and algorithm introduced in this paper using an optimization submitted in a bug report for the Apache Ant build tool [3]. Fig. 1 contains a program function `srch_strm` and an optimization of the function `srch_strm'` submitted in Apache Bug Report #34464 [2]. The actual original and optimized programs submitted in Bug Report #34464 use additional variables and control structure, and were written in Java, but have been simplified to `srch_strm` and `srch_strm'` in Fig. 1 to simplify the discussion. However, we have implemented our algorithm as a tool

that checks the equivalence of the actual programs (we translated the programs to C++ by hand so that we could apply our checker, which uses the LLVM compiler framework [19]).

`srch_strm` and `srch_strm'` implement an equivalent search for a substring in an input stream. Both functions take a string `s`, and read characters from a stream until they have read a string that contains `s`. `srch_strm` implements the search by constructing an empty string `sb` and iteratively checking if `sb` contains `s` as a subsequence. In each iteration, `srch_strm` calls `find(sb, s)`, which returns `True` if and only if `sb` constrains `s` as a subsequence. If `find(sb, s) = True`, then `srch_strm` returns `sb`. Otherwise, `srch_strm` gets a character `c` from the stream, appends `c` to `sb`, and iterates again.

An Apache developer observed that while `srch_strm` outputs the correct value for each input string, it is inefficient due to how it uses the string library functions `find` and `append`. In Apache Bug Report #34464, the developer submitted a patch to `srch_strm`, called `srch_strm'`, that is functionally equivalent to `srch_strm`, but which the developer measured to be more efficient than `srch_strm`. `srch_strm'` is structured similarly to `srch_strm`, but executes more efficiently by only searching for `s` in a sufficiently long suffix of `sb`. In particular, `srch_strm'` maintains an integer variable `pos` that stores the position in `sb` from which `srch_strm'` searches for `s`. In each iteration, `srch_strm'` constructs the substring of `sb` starting at `pos`, `sub(sb, pos)`, tries to find `s` in `sub(sb, pos)`, and if it fails, gets a new character `c` from the stream, appends `c` to `sb`, and iterates again.

An Apache developer submitted `srch_strm'` with an informal argument that it is semantically equivalent to `srch_strm`, but ideally, the developer would submit `srch_strm'` accompanied by a proof that could be checked automatically to determine that `srch_strm` is equivalent to `srch_strm'`. Existing analyses for constructing automatically-checkable proofs of equivalence construct a *simulation relation* from P to P' , which shows that every execution of P corresponds to an execution of P' that returns the same value [20, 23, 25]. A simulation relation \sim from P to P' is a binary relation from the states of P to the states of P' such that: (1) \sim relates each initial state of P to the state in P' with equal values in each variable; (2) \sim relates each return state of P to a state of P' with the same return value; and (3) if \sim relates a state q_0 of P to a state q'_0 of P' and q_0 transitions to a state q_1 of P , then q'_0 transitions, possibly over multiple steps, to some state q'_1 of P' such that \sim relates q_1 to q'_1 .

One simulation relation from `srch_strm` to `srch_strm'`, under an intuitive semantics of the library functions `append`, `find`, `sub`, and `len`, is:

$$(L_0, L'_0) : s = s' \tag{1}$$

$$(L_1, L'_1) : s = s' \wedge sb = sb' \tag{2}$$

$$(L_2, L'_2) : sb = sb' \tag{3}$$

The simulation relation \sim_{ex} of Formulas (1)–(3) is represented as a map from a pair of program labels to a formula that describes pairs of program states. A state q of `srch_strm` at label L is related to a state q' of `srch_strm'` at label L'

if the values of the variables in q and q' satisfy the formula mapped from (L, L') (where the variables of q' are primed). \sim_{ex} relates all initial states of `srch_strm` to initial states of `srch_strm'` with equal values for \mathbf{s} (Formula (1)), all return states of `srch_strm` to return states of `srch_strm'` that return the same value (Formula (3)), and states at the loop head of `srch_strm` to states at the loop head of `srch_strm'` with equal values for \mathbf{s} and \mathbf{sb} (Formula (2)).

\sim_{ex} is a straightforward instance of the definition of a simulation relation for `srch_strm` and `srch_strm'`. However, the fact that \sim_{ex} satisfies condition (2) of a simulation relies on the semantics of the library functions `append`, `find`, and `sub` called by `srch_strm` and `srch_strm'`. Unfortunately, in practice it is difficult to automatically infer accurate specifications for library functions, as such functions may be unavailable or difficult to analyze.

Fortunately, while a programmer may not be able to give a complete formal specification of a library, they often understand a weaker, partial specification that implies the equivalence of a particular optimization. For example, `srch_strm` and `srch_strm'` are equivalent under the assumptions that (1) if the length of string sb is less than the length of string s , then `find(sb, s) = False` and (2) if `find(sb, s) = False`, then \mathbf{s} is a subsequence of the concatenation of \mathbf{s} with a character c if and only if \mathbf{s} is a subsequence of the suffix of \mathbf{s} and c of length equal to the length of \mathbf{s} :

$$\forall sb, s. \text{len}(sb) - \text{len}(s) < 0 \implies \neg \text{find}(sb, s) \quad (4)$$

$$\forall sb, s, c. \neg \text{find}(sb, s) \implies (\text{find}(\text{append}(sb, c), s) \quad (5)$$

$$\iff \text{find}(\text{sub}(sb, \text{len}(sb) - \text{len}(s)), s))$$

Based on the insight that programmers can often reliably validate partial specifications of libraries, in this work we introduce the *abductive equivalence* problem AEQ (§3.2). An abductive equivalence problem is defined by an original program P , optimized program P' , and an *oracle*, which models a programmer, that takes a formula φ describing the library functions called by P and P' and accepts φ if the oracle's model of the library functions satisfies φ . A solution to the problem is a simulation relation from P to P' under the oracle's model of the library functions.

We present a sound algorithm for AEQ, called ChkAEQ (§3.3), that extends existing algorithms for checking program equivalence [23]. Like algorithms for checking equivalence, ChkAEQ first asserts that the return values of input programs P and P' are equal, and then reasons backwards over the executions of P and P' to construct a simulation relation from P to P' represented as a map from pairs of program control labels to formulas in a logic that describes the states of the program. The key feature of ChkAEQ is that as it constructs a relation \sim from the states of P to P' , it applies an *abductive theorem prover* to construct a condition on the library functions (i.e., a *library condition*) that implies that \sim is a simulation relation. If ChkAEQ finds a simulation relation and sufficient library condition, it queries the input oracle on the library condition to determine if the oracle's model satisfies the library condition. If the library oracle validates the condition, then ChkAEQ returns the simulation relation. Otherwise, if the oracle

$$\begin{array}{ll}
\text{block} := L : \text{instr}; \text{term} & L \in \text{Labels} \\
\text{instr} := x_0 := f(x_1, \dots, x_n) \mid x_0 := g(x_1, \dots, x_n) & \{x_i\}_i \subseteq \text{Vars}; f \in \text{Ops}; g \in \text{LibOps} \\
\text{term} := \text{return } x \mid \text{br } x ? L_t : L_f & L_t, L_f \in \text{Labels}; x \in \text{Vars}
\end{array}$$

Fig. 2. Syntax of the programming language IMP, described in §3.1. An IMP program is a set of blocks.

refutes the condition, then ChkAEQ uses the refutation to continue to search for a simulation relation.

For `srch_strm` and `srch_strm'` (Fig. 1), ChkAEQ could infer that if each string s can be found in each string sb (i.e., $\forall sb, s. \text{find}(sb, s)$ (6) is valid), then \sim_{ex} is a simulation relation. However, a programmer serving as a library oracle would refute Formula (6). ChkAEQ would then use the refutation to search for a condition consistent with the negation of Formula (6), and would eventually find the library conditions of Formulas (4) and (5).

3 Abductive Equivalence

In this section, we formally define the abductive equivalence problem and algorithm. In §3.1, we define the syntax and semantics of a simple imperative language IMP. In §3.2, we define the abductive equivalence problem for IMP programs. In §3.3, we present a sound algorithm for solving the abductive equivalence problem.

3.1 IMP: a Simple Imperative Language

IMP Syntax. An IMP program updates its state by executing a sequence of program and library operations. The IMP language (Fig. 2) is defined over a set of variable symbols `Vars`, a set of control labels `Labels`, a set of language function symbols `Ops`, and a set of library function symbols `LibOps`, where `Vars`, `Labels`, `Ops`, and `LibOps` are mutually disjoint. `Labels` contains a label `RET` that does not label any block of an IMP program (`RET` is used to define the semantics of a return instruction; see “IMP Semantics”).

An IMP program is a set of basic blocks. Each IMP program P contains one initial block labeled with a $\iota(P) \in \text{Labels}$. Each basic block `block` is a label followed by an instruction and block terminator. An instruction is an assignment of either a language operation or a library operation. A block terminator is either a return instruction or a conditional branch.

IMP Semantics. The operational semantics of IMP (Fig. 3) defines how a basic block transforms a given state under a given model of library operations. An IMP *state* is a label paired with a *valuation*, which is a map from each

$$\begin{array}{l}
\sigma_b^m \llbracket L : \text{instr}; \text{term} \rrbracket (L_0, V) \equiv \text{if } L_0 = L \text{ then } \sigma_t \llbracket \text{term} \rrbracket (\sigma_i^m \llbracket \text{instr} \rrbracket (V)) \text{ else } \perp \quad (6) \\
\hline
\sigma_i^m \llbracket \mathbf{x}_0 := \mathbf{f}(\mathbf{x}_1) \rrbracket (V) \equiv V[x_0 \mapsto \text{langmodel}(\mathbf{f})(V(x_1), \dots, V(x_n))] \quad (7) \\
\sigma_i^m \llbracket \mathbf{x}_0 := \mathbf{g}(\mathbf{x}_1) \rrbracket (V) \equiv V[x_0 \mapsto m(\mathbf{g})(V(x_1), \dots, V(x_n))] \quad (8) \\
\hline
\sigma_t \llbracket \text{return } \mathbf{x} \rrbracket (V) \equiv (\text{RET}, V[\text{rv} \mapsto V(x)]) \quad (9) \\
\sigma_t \llbracket \text{br } \mathbf{x} ? L_t : L_f \rrbracket (V) \equiv ((\text{if } V(x) \neq 0 \text{ then } L_t \text{ else } L_f), V) \quad (10)
\end{array}$$

Fig. 3. Operational semantics of IMP. σ_b^m is the operational semantics of a block and σ_i^m is the operational semantics of an instruction under library model m . σ_t is the operational semantics of a block terminator. In σ_b^m , \perp denotes the undefined value. In σ_i^m and σ_t , for $a, b \in \mathbb{Z}$, $V[a \mapsto b]$ maps a to b , and maps $c \neq a$ to $V(c)$. In σ_t , $\text{rv} \in \text{Vars}$ stores the return value of the program.

program variable to an integer value: $\text{Valuations} = \text{Vars} \rightarrow \mathbb{Z}$ and $\text{States} = \text{Labels} \times \text{Valuations}$. A *library model* $m : \text{LibOps} \rightarrow (\mathbb{Z}^* \rightarrow \mathbb{Z})$ maps each library function to a function from a vector of integers to an integer.

The semantic function of a block σ_b^m (Fig. 3, Eqn. (6)) defines how a basic block transforms a state under m . The semantic function of an instruction σ_i^m (Fig. 3, Eqns. (7) and (8)) defines how an instruction updates a valuation. If an instruction assigns the result of a language operation, then the value of the operation is defined by a fixed *language model* $\text{langmodel} : \text{Ops} \rightarrow (\mathbb{Z}^* \rightarrow \mathbb{Z})$ (Fig. 3, Eqn. (7)). If an instruction assigns the result of a library operation, then the value of the operations is defined by m (Fig. 3, Eqn. (8)). The semantic function of a block terminator σ_t defines how a block terminator transforms a state (Fig. 3, Eqns. (9) and (10)).

3.2 The Abductive Equivalence Problem

To define the abductive equivalence problem, we adopt the definition of a simulation relation [20], which has been used to define the classical equivalence problem [23, 25]. A simulation relation from an IMP program P to an IMP program P' is a relation from states of P to states in P' that implies that if from inputs I , P has an execution that returns value v , then from I , P' also has an execution that returns v .

Defn. 1 Let IMP program P be compatible with IMP program P' if every variable of P corresponds to a variable of P' : $\text{Vars}(P) \subseteq \text{Vars}(P')$. A simulation relation from P to a compatible program P' under library model $m : \text{LibOps} \rightarrow (\mathbb{Z}^* \rightarrow \mathbb{Z})$ is a relation $\sim \subseteq \text{States} \times \text{States}$ from the states of P to the states of P' such that:

1. Initial states: \sim relates each initial state of P to its analogous state in P' . For valuations $V, V' \in \text{Valuations}$, if for each $x \in \text{Vars}(P)$, $V(x) = V'(x)$, then $\sim ((\iota(P), V), (\iota(P'), V'))$.
2. Return states: \sim relates each return state of P to a return state of P' that returns an equal value. For the distinguished variable $\text{rv} \in \text{Vars}$ holding the

Algorithm 1 An abductive equivalence algorithm ChkAEQ. Takes as input a \mathcal{T}_{lib} oracle \mathcal{O}_m and two IMP programs P and P' , and constructs a solution to $\text{AEQ}(m, P, P')$ (Defn. 2). $\text{Interact}(A)$ returns a simulation relation under a library condition that implies A and is validated by \mathcal{O}_m . ChkAEQ is discussed in §3.3.

```

1: function ChkAEQ( $\mathcal{O}_m, P, P'$ )
2:   function Interact( $A$ )
3:      $(C, \sim) := \text{SimRel}(P, P', A)$ 
4:     if  $\mathcal{O}_m(C)$  then return  $\sim$ 
5:     else return Interact( $A \wedge \neg C$ )
6:   end if
7: end function
8: return Interact(True)
9: end function

```

return value of the function, and for each $V \in \text{Valuations}$, there is some $V' \in \text{Valuations}$ such that $V(\text{rv}) = V'(\text{rv})$ and $\sim((\text{RET}, V), (\text{RET}, V'))$.

3. Consecution: if \sim relates a state s_0 of P to a state s'_0 of P' and s_0 transitions to a state s_1 of P , then s'_0 eventually transitions to a state s'_1 of P' such that \sim relates s_1 to s'_1 . For program P , let the transition relation $\rightarrow_P \subseteq \text{States} \times \text{States}$ relate states connected by the transition function of P : for states $s_0, s_1 \in \text{States}$, $\rightarrow_P(s_0, s_1)$ if and only if there is some block $B \in P$ such that $s_1 = \sigma_b^m \llbracket B \rrbracket(s_0)$. Let $\rightarrow_P^* \subseteq \text{States} \times \text{States}$ be the reflexive transitive closure of \rightarrow_P . For $s_0, s'_0, s_1 \in \text{States}$, if $\sim(s_0, s'_0)$ and $\rightarrow_P(s_0, s_1)$, then there is some $s'_1 \in \text{States}$ such that $\rightarrow_{P'}^*(s'_0, s'_1)$ and $\sim(s_1, s'_1)$.

The abductive equivalence problem for programs P and P' and library model m is to find a simulation relation from P to P' under m using P, P' , and an oracle for m that answers Boolean queries about properties of m . Intuitively, the oracle formalizes the role of a programmer who can answer queries about the specification of a library. The oracle answers queries on properties expressed as formulas of a logical theory \mathcal{T}_{lib} whose models describe IMP's library operations. For each library operation $g \in \text{LibOps}$, there is an uninterpreted function symbol g in \mathcal{T}_{lib} . The only predicate of \mathcal{T}_{lib} is the equality relation $=$. The set of library conditions $\text{Forms}(\mathcal{T}_{lib})$ are the first-order formulas of \mathcal{T}_{lib} , and for library condition $\varphi \in \text{Forms}(\mathcal{T}_{lib})$, $m \models \varphi$ denotes that m is a model of φ . For library model m , the oracle $\mathcal{O}_m \subseteq \text{Forms}(\mathcal{T}_{lib})$ accepts a \mathcal{T}_{lib} formula if and only if m satisfies the formula: for $\varphi \in \text{Forms}(\mathcal{T}_{lib})$, $\mathcal{O}_m(\varphi)$ if and only if $m \models \varphi$.

Defn. 2 For library model m and IMP programs P and P' , the abductive equivalence problem $\text{AEQ}(m, P, P')$ is, given \mathcal{O}_m, P , and P' , to find a simulation relation \sim from P to P' under m .

3.3 A Sound Algorithm for Abductive Equivalence

Interacting with an Oracle to Solve AEQ. In this section, we present an algorithm that soundly tries to solve the abductive equivalence problem (Defn. 2),

Algorithm 2 *SimRel*: takes an original program P , optimized program P' , and library condition A , and constructs a library condition C such that $C \implies A$ and a simulation relation from P to P' under C . sym_{ret} , ConsLocRel , and ConsecSimRel are discussed in §3.3.

```

1: function SimRel( $P, P', A$ )
2:    $\text{lr} \leftarrow \text{ConsLocRel}(P, P')$ 
3:   function ConsecSimRel( $C, \text{sym}, W$ )
4:     if  $W = \emptyset$  then return ( $C, \text{sym}$ )
5:     else
6:        $(L, L') \leftarrow \text{RemElt}(W)$ 
7:        $\text{cs} := \bigvee \{ \text{wp}_P^{P'} \llbracket (L, L'), (L_1, L'_1) \rrbracket (\text{sym}(L_1, L'_1)) \mid (L_1, L'_1) \in \text{succ}_r(L, L') \}$ 
8:        $\text{suff} := \text{ATP}(\text{sym}(L, L') \implies \text{cs})$ 
9:       if  $\text{IsTConsistent}(\text{suff})$  then  $C := C \wedge \text{suff}$ 
10:      else
11:         $\text{sym}(L, L') := \text{cs}$ 
12:         $W := W \cup \text{preds}_r(L, L')$ 
13:      end if
14:      return ConsecSimRel( $C, \text{sym}, W$ )
15:    end if
16:  end function
17:   $(A_{cs}, \text{sym}_{cs}) := \text{ConsecSimRel}(A, \text{sym}_{\text{ret}}, \text{dom}(\text{lr}) \cup \text{img}(\text{lr}))$ 
18:  if  $\text{IsTValid}(\text{initeq} \implies \text{sym}_{cs}(\iota(P), \iota(P')))$  then return ( $A_{cs}, \text{sym}_{cs}$ )
19:  else abort
20:  end if
21: end function

```

called *ChkAEQ* (Alg. 1). *ChkAEQ* tries to solve an abductive equivalence problem $\text{AEQ}(m, P, P')$ as follows. *ChkAEQ* first defines a function *Interact* (Alg. 1, lines [2]–[6]) that takes a library condition A , and constructs a simulation relation from P to P' under m if $m \models A$. When *ChkAEQ* is successful, it returns the simulation relation constructed by applying *Interact* to *True* (Alg. 1, line [8]). However, *ChkAEQ* may fail or not terminate.

Interact constructs a simulation relation by applying a function *SimRel* (Alg. 1, line [3]) that takes an original program P , optimized program P' , and a library condition A , and constructs (1) a library condition C such that $C \implies A$ and (2) a simulation relation from P to P' under each library model that satisfies C (in which case, we say that \sim is a simulation from P to P' *under* C). If \mathcal{O}_m accepts C (Alg. 1, line [4]), then *Interact* returns \sim (Alg. 1, line [4]). Otherwise, *Interact* calls itself with a stronger library condition that asserts that C is not valid (Alg. 1, line [5]).

Constructing a Simulation Relation. *SimRel* constructs a simulation relation represented as a symbolic state relation, which is a function from pairs of labels to a formula of a theory \mathcal{T} for which each model defines a library model, a state of P , and a state of P' . The theory \mathcal{T}_{IMP} describes states of fixed programs P and P' . For each $\mathbf{f} \in \text{Ops}$, let there be a unary function f in the logical theory

\mathcal{T}_{IMP} . For each program variable $x \in \text{Vars}(P)$, let there be a \mathcal{T}_{IMP} constant x , and for each $x \in \text{Vars}(P')$, let there be a \mathcal{T}_{IMP} constant x' that does not correspond to any variable in $\text{Vars}(P)$. Let the only predicate of \mathcal{T}_{IMP} be the equality predicate, which \mathcal{T}_{IMP} shares with \mathcal{T}_{lib} (§3.2). Let the combination of \mathcal{T}_{lib} and \mathcal{T}_{IMP} be $\mathcal{T} = \mathcal{T}_{\text{lib}} + \mathcal{T}_{\text{IMP}}$ [24]. A symbolic state relation $\text{sym} : (\text{Labels} \times \text{Labels}) \rightarrow \text{Forms}(\mathcal{T})$ relates states $(L, v), (L', v') \in \text{States}$ under library condition C if for each library model $m \models C$, $m \cup v \cup v' \models \text{sym}(L, L')$.

SimRel constructs a library condition A_{cs} such that $A_{cs} \implies A$, and a symbolic state relation sym_{cs} that satisfies the consecution condition of a simulation relation (Defn. 1, item 3) under A_{cs} . SimRel defines a *label transition relation* $\text{lr} \subseteq (\text{Labels} \times \text{Labels}) \times (\text{Labels} \times \text{Labels})$, which defines the domain of sym_{cs} , by applying a function ConsLocRel (Alg. 2, line [2]). ConsLocRel can be implemented using known techniques and heuristics from classical equivalence checking [23]. SimRel then defines a function ConsecSimRel (Alg. 2, lines [3]–[16]) that takes as input (1) a library condition C , (2) a symbolic state relation sym , and (3) a workset W of label pairs on which sym may not satisfy Defn. 1, item 3 under C , and constructs a library condition A_{cs} and a state relation sym_{cs} that satisfy the consecution condition, and such that A_{cs} implies C . To construct sym_{cs} and A_{cs} , SimRel applies ConsecSimRel to A , a simulation relation sym_{ret} , and all pairs of labels in the domain and image of lr (Alg. 2, line [17]). sym_{ret} satisfies the return condition of a simulation relation, defined in Defn. 1, item 1.

SimRel then checks that sym_{cs} satisfies the condition for a simulation relation on initial blocks (Defn. 1, item 1) by checking that $\text{initeq} \equiv \bigwedge_{x \in \text{Vars}(P)} x = x'$ implies the relation of states at the initial blocks of P and P' (Alg. 2, line [18]). If so, then SimRel returns $(A_{cs}, \text{sym}_{cs})$ as a simulation relation (Alg. 2, line [18]). Otherwise, SimRel fails (Alg. 2, line [19]).

ConsecSimRel (Alg. 2, lines [3]–[16]) first checks if its workset of labels W is empty (Alg. 2, line [4]), and if so, returns its input library condition C and input state relation sym (Alg. 2, line [4]). Otherwise, ConsecSimRel chooses a pair of labels (L, L') from W (Alg. 2, line [6]) on which it will update sym . For $\varphi \in \text{Forms}(\mathcal{T})$, let $\text{wp}_P^{P'} \llbracket (L_0, L'_0), (L_1, L'_1) \rrbracket (\varphi)$ be the formula whose models define states that transition to states in φ over steps of execution in P from L_0 to L_1 and steps of execution in P' from L'_0 to L'_1 ($\text{wp}_P^{P'}$ is defined from the semantics of IMP (Fig. 3) using well-known techniques [11]). ConsecSimRel constructs cs , the disjunction of the weakest precondition of each formula to which sym maps each successor of L and L' under lr (Alg. 2, line [7]). ConsecSimRel then tries to construct a library condition $\text{suff} \in \text{Forms}(\mathcal{T}_{\text{lib}})$ that is consistent and implies that $\text{sym}(L, L')$ implies cs . To construct suff , ConsecSimRel applies an abductive theorem prover ATP (Alg. 2, line [8]; ATP is discussed in [12], App. A). If suff is consistent, then ConsecSimRel conjoins suff to C (Alg. 2, line [9]). Otherwise, ConsecSimRel updates sym to map L and L' to cs , and adds each predecessor of L and L' under lr to W (Alg. 2, lines [11]–[12]). ConsecSimRel then calls itself recursively on its updated library condition, symbolic state relation, and workset.

Correctness of ChkAEQ. AEQ is at least as hard as determining if two IMP programs are equivalent. IMP is Turing-complete, and thus AEQ is undecidable. However, ChkAEQ is sound, but not complete, for AEQ. ChkAEQ also does not pose redundant queries to its oracle. These claims are formalized in an extended version of this paper ([12], Sec. 3).

Suppose that programs P and P' are not equivalent under library model m . Then for $\text{AEQ}(mP, P')$, ChkAEQ either will not terminate, or will abort when it fails to find a simulation relation that relates the program points guessed by ConsLocRel. It would be interesting to extend ChkAEQ so that it simultaneously searches for sufficient library conditions under which programs are equivalent, or sufficient library conditions that prove that the programs are definitely *not* equivalent. If a programmer submits a patch that is not equivalent to their original program, then ChkAEQ extended in this way could explain to the programmer why their patch is incorrect.

4 Experiments

We carried out a set of experiments to determine if programmers can apply ChkAEQ (§3) to validate practical optimizations. The experiments were designed to answer the following questions:

1. Given a function from a real-world program and its optimization, can ChkAEQ quickly find a library condition that is sufficient to prove that the programs are equivalent?
2. Can ChkAEQ find library conditions that are small and easy for a programmer to validate?

To answer these questions, we implemented ChkAEQ as a tool, `chklibs`, and applied `chklibs` to a set of program functions and their optimizations. Each function was taken from a mature, heavily-used program, namely the Apache software suite, Mozilla Suite, or MySQL database. Each original program function was the subject of a bug report reporting that the function’s behavior was correct, but that its performance was inefficient. Each corresponding optimized function was the patched, optimized function provided in the bug report. We interacted with `chklibs` to find library conditions that were sufficient to prove that the optimization was correct, and were valid according to our understanding of the libraries.

The results of the experiments indicate that ChkAEQ can be applied to validate practical optimizations. In particular:

1. `chklibs` quickly inferred library conditions that were sufficient to prove equivalence. `chklibs` usually found validated sufficient library conditions in less than a second, and always found validated sufficient library conditions in less than 30 seconds (see Table 1).
2. `chklibs` often inferred syntactically compact sufficient library conditions. `chklibs` usually needed to suggest less than 10 disjunctive clauses until it

Benchmark Data					chklibs Performance			
Program Name	Bug ID	LoC			Analysis Time (s)	Num. Clauses	Avg. Query Size	
		Org.	Opt.	Diff.			Clause Size	Num. Preds
Apache	19101	27	28	5	0.325	2	9.5	6.0
	34464	23	20	34	18.188	5	9.0	6.4
	44408	51	52	6	0.050	1	8.0	6.0
	45464	569	570	6	0.165	1	8.0	6.0
	48778	30	28	16	0.534	7	11.4	6.0
Mozilla	103330	217	216	5	0.064	2	18.0	6.0
	124686	198	198	4	0.096	1	278.0	6.0
	267506	182	184	9	0.507	5	8.0	6.0
	409961	54	57	12	0.795	3	47.0	6.0
MySQL	38769	223	227	4	0.169	2	11.0	6.0
	38824	346	321	18	29.894	13	179.2	6.0

Table 1. Experimental data from using `chklibs`. The data given for each benchmarks program includes the name and of the source program, the bug report that presented the optimization, the number of lines of code of the original and optimized program functions, and the number of lines output by `diff` on the original and optimized programs. The data measuring `chklibs`’s performance includes the time taken by `chklibs` to construct a simulation relation (in seconds), the number of clauses on which `chklibs` queried the user, the average size of (i.e, number of logical symbols in) the clause, and the average number of predicates in each clause.

suggested a sufficient set of clauses. The clauses always contained less than 10 predicates (and usually contained less than 5 predicates), and with some exceptions discussed below, were small enough that a programmer should be able to reason about their validity.

In §4.1, we describe in detail our procedure for evaluating `ChkAEQ`. In §4.2, we present and analyze the results of applying `ChkAEQ`.

4.1 Experimental Procedure

Implementation. `chklibs` solves the abductive equivalence problem for the LLVM [19] intermediate language. To implement `chklibs`, we extended the operational semantics of IMP (§3.1) to an operational semantics for the LLVM intermediate language, which included describing various language features such as structs and pointers. Such an extension is standard, and we omit its details. `chklibs` is implemented in about 5,000 lines of OCaml code, and uses the Z3 theorem prover [7] to implement the abductive theorem prover ATP ([12], App. A). `chklibs` simplifies each query and presents the query to the user as a conjunction of disjunctive clauses. We discuss simplifications that `chklibs` applies to queries in an extended version of this paper ([12], App. B).

Evaluation. To evaluate `chklibs`, we used it to validate a set of optimizations submitted to improve the performance of real-world applications. In particular,

we collected a set of bug reports from the public bug databases of Apache software suite [1], Mozilla Suite [21], and MySQL database [22] that each reported a performance issue and included a patch to fix the issue. We compiled each original program and its patch to the LLVM intermediate language. If a program and patch were originally implemented in a language supported by LLVM, such as C or C++, then we compiled the programs by applying the appropriate LLVM compiler front-end (`clang` or `clang++`, respectively). Otherwise, we rewrote the program functions by hand in C source code and compiled the source to the LLVM intermediate language by applying `clang`.

We applied `chklibs` to each original and optimized program, and interacted with `chklibs` to find library conditions that `chklibs` determined to be sufficient, and which were valid under our understanding of the libraries and program functions described informally in the bug report. In other words, we served as the library oracle introduced in §3.3. For each benchmark, we observed whether or not `chklibs` found library conditions that we believed to be valid, measured the total time spent by `chklibs` to infer sufficient specifications, measured the number of queries issued by `chklibs`, and measured the size of each query.

4.2 Results and Analysis

Results. Table 1 contains the results of applying `chklibs`. Each row in Table 1 contains data for a benchmark original and optimized program. In particular, Table 1 contains the name of the program from which the benchmark was taken, the ID of the bug report in which the optimization was submitted, the number of lines of code of the original and optimized program functions, the number of lines output by `diff` on the original and optimized programs, the time spent by `chklibs` to construct a validated simulation relation (not including the time spent by us to validate or refute a query posed by `chklibs`), the number of clauses on which `chklibs` queried the user, and average size of (i.e., the number of all logical symbols in) the clauses, and the average number of predicates in each clause. The size of the clause is the number of logical symbols in the clause.

Analysis. The data presented in Table 1 indicates that `chklibs` can be applied to suggest sufficient library conditions for equivalence that can often be easily validated by a programmer. In benchmarks where `chklibs` took an unusually long time to find validated sufficient conditions (Apache Bug 34464 and MySQL Bug 38824), `chklibs` posed a proportionally large number of queries that we refuted. In benchmarks where `chklibs` queried the user on an unusually large set of clauses (Apache Bug 48778 and MySQL Bug 38824), the original and optimized programs called different library functions at a proportionately large set of callsites. In benchmarks where `chklibs` queried the user on unusually large formulas (Mozilla Bug 124686 and MySQL Bug 38824), the formulas typically were constructed from equality conditions over addresses in the program that seem unlikely to alias. We believe that the size of these queries could be reduced drastically by combining `chklibs` with a more sophisticated alias analysis, or

a programmer with a more detailed understanding of the calling conventions of the original and optimized function.

We have provided a website¹ that contains, for each benchmark, the list of all queries generated by `chklibs` that we answered. Apache Bug #19101 and Mozilla Bug #409961 illustrate a common kind of optimization in which a programmer performs an interprocedural version of a classic compiler optimization: in Apache Bug #19101 and Mozilla Bug #409961, the optimization is loop-invariant code motion. To validate the optimizations, `chklibs` generates queries that determine sufficient library conditions to support the optimization. For Apache Bug #19101, `chklibs` queries if a method called within a loop does not change the values stored at particular fields of the calling object. For Mozilla Bug #409961, `chklibs` must determine that the final values returned by the programs are equivalent, even though some intermediate values computed within a loop may be different as a result of the optimization.

Apache Bug #48778 illustrates that in practical optimizations, there may be multiple consistent conditions on library that are sufficient to support an optimization. In Apache Bug #48778, the values returned by library functions determine the values of control-flow guards. `chklibs` correctly determines that if the guard values are constant, then the programs are equivalent. We refuted the corresponding library condition, which caused `chklibs` to eventually find a valid library condition equating the return values of particular library functions.

MySQL Bug #38769 illustrates how `chklibs` can make explicit supporting conditions that may be non-obvious to the developers. MySQL Bug #38769 optimizes a loop over an array by replacing a constant loop bound in the original program with the result of a method call, which may be a lower value. `chklibs` first determines that if the result of the method call is equal to the replaced constant, then the programs are equal. We refuted this condition, at which point `chklibs` determined that if all entries after the bound returned by the call are null, then the programs are equivalent, and we accepted this condition. For the report in which the patch was submitted, a developer notes that this condition did not hold for older versions of MySQL.

5 Related Work

Previous work [15] identified *performance bugs*, i.e., functionally correct but inefficient code, as a serious problem in commonly-used applications. In that work, the authors studied a set of performance bugs for five software suites, namely Apache, Chrome, GCC, Mozilla, and MySQL, derived a set of rules for identifying performance bugs manually from performance-bug reports, and statically checked programs to find new performance bugs that satisfy the rules. We have presented a technique and tool that allows a developer who submits a patch of a performance bug to validate conditions under which the patch preserves functionality of the program. We have applied to the tool to patches of bugs identified in the previous work on performance bugs.

¹ <http://pages.cs.wisc.edu/~wrharris/chklibs/>

Much existing work has focused on determining the equivalence of programs. Translation validation [23, 25] is the problem of determining if a source program is equivalent to an optimized program, and is often applied to validate the correctness of the phases of an optimizing compiler. Regression verification [10] determines if a program and a similar revision of the program are equivalent. Semantic differencing [14] summarizes the different behaviors of two programs. Symbolic execution has been applied to determine the equivalence of loop-free programs [5]. In this work, we address the problem of taking an original and optimized program and inferring conditions on the libraries invoked by the programs that are sufficient to prove that the programs are equivalent, and that are validated by an oracle who understands the libraries. To solve the problem, we have extended an existing analysis for checking the equivalence of imperative programs with loops [23] to use the results of an abductive theorem prover.

The SymDiff project [16–18] shares our goal of determining under what conditions two programs are correct. Existing work in SymDiff takes a concurrent program, constructs a sequential version of the program, and treats the sequential version as a reference implementation, searching the concurrent program only for bugs triggered by inputs that cause no error in the sequential program [16]. Existing work on conditional equivalence [17] takes an original and optimized program and infers sufficient conditions on the inputs of a program under which the original and optimized programs are equivalent, where the space of conditions forms a lattice. In contrast, our work interacts with a user to infer sufficient conditions on the libraries invoked by an original and optimized program, and represents conditions as logical formulas. Given that the spaces of conditions described in techniques based on conditional equivalence must form a lattice, it is not immediately clear how to extend such a technique to interact with a user who may refute an initial condition suggested by the technique.

Recent work has extended the problem of deciding if a program always satisfies an assertion to an abductive setting, in which the problem is to find assumptions on the state of a program that imply that the program satisfies an assertion, and are validated by an oracle that answers queries about program states [8]. That work presents an algorithm that constructs sufficient assumptions by finding a minimum satisfying assignment of variables in a given formula [9], universally quantifying all unassigned variables, and eliminating the quantified variables using symbolic reasoning. Our work extends a different traditional problem in program analysis, that of checking program equivalence, to an abductive setting. While work on abductive assertion checking assumes that the theories for describing states of a program support quantifier elimination (e.g., linear arithmetic), we consider inferring assumptions for theories that may describe arbitrary library functions. Accordingly, our analysis applies an abductive theorem prover that does not assume that the theories modeling the semantics of a program support quantifier elimination, and instead generates assumptions as Boolean combinations of equality predicates.

Our approach to the abductive equivalence problem collects information from a programmer about the libraries that a program uses by querying the program-

mer for the validity of purely equational formulas over the library functions, and propagates the logical consequence of the equalities to the rest of the analysis. Our approach is inspired by equality propagation [24], which is a technique for combining solvers for theories whose only shared predicate is equality to solve a formula defined in the combination of the theories. Essentially, our approach uses the programmer as a theory solver for the theory modeling library functions.

References

1. Apache, Jan. 2013. <http://apache.org>.
2. Apache bug #34464, Jan. 2013. http://issues.apache.org/bugzilla/show_bug.cgi?id=34464.
3. Apache Ant, May 2012. <http://ant.apache.org>.
4. T. Ball and S. K. Rajamani. The SLAM project: debugging system software via static analysis. In *POPL*, 2002.
5. E. M. Clarke, D. Kroening, and K. Yorav. Behavioral consistency of C and Verilog programs using bounded model checking. In *DAC*, 2003.
6. P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, 1977.
7. L. M. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, 2008.
8. I. Dillig, T. Dillig, and A. Aiken. Automated error diagnosis using abductive inference. In *PLDI*, 2012.
9. I. Dillig, T. Dillig, K. L. McMillan, and A. Aiken. Minimum satisfying assignments for SMT. In *CAV*, 2012.
10. B. Godlin and O. Strichman. Regression verification. In *DAC*, 2009.
11. D. Gries. *The Science of Programming*. Springer, 1981.
12. W. R. Harris, G. Jin, S. Lu, and S. Jha. Validating library usage interactively, Jan. 2013. http://pages.cs.wisc.edu/~wrharris/validating_library_usage.pdf.
13. T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Lazy abstraction. In *POPL*, 2002.
14. D. Jackson and D. A. Ladd. Semantic diff: A tool for summarizing the effects of modifications. In *ICSM*, 1994.
15. G. Jin, L. Song, X. Shi, J. Scherpelz, and S. Lu. Understanding and detecting real-world performance bugs. In *PLDI*, 2012.
16. S. Joshi, S. K. Lahiri, and A. Lal. Underspecified harnesses and interleaved bugs. In *POPL*, 2012.
17. M. Kawaguchi, S. K. Lahiri, and H. Rebelo. Conditional equivalence. Technical Report MSR-TR-2010-119, Microsoft Research, Oct 2010.
18. S. K. Lahiri, C. Hawblitzel, M. Kawaguchi, and H. Rebêlo. SymDiff: A language-agnostic semantic diff tool for imperative programs. In *CAV*, 2012.
19. The LLVM compiler infrastructure.
20. R. Milner. *Communication and concurrency*. PHI Series in computer science. Prentice Hall, 1989.
21. Mozilla – home of the Mozilla Project, 2011. <http://www.mozilla.org/>.
22. MySQL: The world’s most popular open source database, 2012. <http://www.mysql.com/>.
23. G. C. Necula. Translation validation for an optimizing compiler. In *PLDI*, 2000.
24. G. Nelson and D. C. Oppen. Simplification by cooperating decision procedures. *ACM Trans. Program. Lang. Syst.*, 1(2):245–257, 1979.
25. A. Pnueli, M. Siegel, and E. Singerman. Translation validation. In *TACAS*, 1998.