

Regression Verification: Proving the Equivalence of Similar Programs

Benny Godlin¹, Ofer Strichman²

¹ CS, Technion, Haifa, Israel.

² Information Systems Engineering, IE, Technion, Haifa, Israel

SUMMARY

Proving the equivalence of successive, closely related versions of a program has the potential of being easier in practice than functional verification, although both problems are undecidable. There are three main reasons for this claim: 1) it circumvents the problem of specifying what the program should do, 2) the problem can be naturally decomposed and hence is computationally easier, and 3) there is an automatic invariant that enables to prove equivalence of loops and recursive functions in most practical cases. Theoretical and practical aspects of this problem are considered. Copyright © 0000 John Wiley & Sons, Ltd.

Received . . .

KEY WORDS: Software verification; Equivalence checking

1. INTRODUCTION

Proving the equivalence of successive, closely related versions of a program has the potential of being easier in practice than applying functional verification to the newer version against a user-defined, high-level specification. There are three reasons for this claim. First, it mostly circumvents the problem of specifying what the program should do. The user can take a no-action ‘default specification’ by which the outputs of the program should remain unchanged if the two programs are run with the same inputs. Second, as shown in this article, there are various opportunities for abstraction and decomposition that are only relevant to the problem of proving equivalence between similar programs, and these techniques make the computational burden less of a problem. Finally, loops and recursion typically do not pose a problem for a fully-automatic proof. The reason for this, intuitively, is that equivalence itself is typically an inductive invariant. In other words, if the recursive calls of two recursive functions return the same values given the same inputs at depth i for some $i > 0$, typically the same holds for recursion depth $i - 1$.

Both functional verification and program equivalence of general programs are undecidable. Coping with the former was declared in 2003 by Tony Hoare as a “grand challenge” to the computer science community [1]. Program equivalence can be thought of as a grand challenge in its own right, but there are reasons to believe, as indicated above, that it is a ‘lower hanging fruit’. The observation that equivalence is easier to establish than functional correctness is supported by past experience with two prominent technologies: *regression testing* [2, 3] – the most popular automated testing technique for software – and *equivalence checking* – the most popular formal verification technique for hardware. In both cases the reference is a previous version of the system. Equivalence

*Correspondence to: Email: ofers@ie.technion.ac.il

checking also demonstrates the difference in the computational effort: it is computationally easier than model-checking, at least under the same assumption made here, namely that the two compared systems are mostly similar. One may argue, however, that the notion of correctness is weaker: rather than proving that a model is ‘correct’, it is proven that it is ‘as correct’ as the previous version. In contrast one may argue that it can still expose functional errors since failing to comply with the equivalence specification indicates that something is wrong with the assumptions of the user. In any case, it might be a feasible venue even in cases where the alternative of functional verification is not.

The problem of proving the equivalence of closely related programs was coined *regression verification* by the authors in a position paper related to the above-mentioned grand challenge [4]. Solving it can be useful wherever regression testing is useful, and in particular for checking backward compatibility. A special case is when a programmer introduces a new performance optimizations or applies *refactoring* [5], which is not supposed to have an impact on the external behavior of the program. It can also be used for *change impact analysis*, i.e., the task of identifying the potential consequences of a change in order to know which functions should be retested.

The idea of proving equivalence between programs is not new, and in fact preceded the idea of functional verification.[†] A detailed survey of earlier work will be given later in this section, but it should be mentioned that as far as the authors are aware no one has focused so far on coping with this problem for general programs in real programming languages nor on how to exploit the fact that large parts of the code in the two compared programs are identical. Ideally the complexity of the solution should be dominated by the (semantic) difference between the two compared programs rather than on their size.

A precise definition of the problem addressed here is the following:

Definition 1 (Regression verification)

Given two programs with corresponding sets of functions $\{f_1, \dots, f_n\}$ and $\{f'_1, \dots, f'_m\}$, a (possibly partial) mapping map_f between them and a definition of function equivalence, *Regression verification* is the problem of finding the set

$$\{(f_i, f'_j) \mid (f_i, f'_j) \in map_f, f_i \text{ and } f'_j \text{ are equivalent}\} .$$

There are many different ways to define the notion of Input/Output equivalence (six different notions of equivalence were defined by the authors in a previous publication [7]). In this paper the focus is on the following definition:

Definition 2 (Partial equivalence of functions)

Two functions[‡] f and f' are said to be *partially equivalent* (p-equiv for short) if any two terminating executions of f and f' starting from the same inputs, emit the same outputs.

The problem of function (or program) equivalence according to this definition can be reduced to one of functional verification of a single program P rather easily: P should simply call the two compared functions consecutively with nondeterministic but equal inputs, and assert that they return the same output. The problem with this direct approach is that it makes no use of the expected similarity of the code. Rather, it solves a monolithic functional verification problem without decomposition. As shown in this article, the similar code structure can be beneficial exactly for this reason.

The rest of the article is structured as follows. Sect. 2 describes briefly an inference rule for proving the partial equivalence of recursive programs. This rule is obviously not complete, but turns out to be strong enough for proving partial equivalence in many realistic cases. In Sect. 3 an algorithm for decomposing the equivalence proof – ideally to the granularity of pairs of functions

[†]In his 1969 paper about axiomatic basis for computer programming [6], Hoare points to previous works from the late 50's on axiomatic treatment of the problem of proving equivalence between programs.

[‡]The term ‘function’ is used here in the same way that it is used in most programming languages, which means that it can have multiple outputs and not necessarily terminate. A C function, for example, can have as outputs a return value, global variables to which it writes, and variables that are passed to it by reference.

– will be presented. Some experiments with the Regression Verification Tool (RVT) that was developed by the authors and with Microsoft’s SymDiff [8][§] are reported in Sect. 4, including a small case study that demonstrates the output of RVT.

Many details about RVT are left out, as well as more references to earlier works. The interested reader may find them in the first author’s thesis [9]. The theoretical background on the inference rule that is used, including detailed semantics of the assumed programming language, can be found in a previous publication by the authors [7]. The emphasis here is on the new decomposition algorithm and on the tool RVT.

Related work This article extends an earlier proceedings version [10] in several ways. Most importantly, it generalizes it to programs with mutual recursion, both in the description of the inference rule and in the decomposition algorithm. It also adds a proof of correctness for the main decomposition algorithm. The inference rules that are used, which were described and proved formally elsewhere [7], are described here for the first time in a way that emphasizes their resemblance to Hoare’s rule for recursive functions. This description seems clearer. Finally, a case study as mentioned above is included.

As mentioned earlier, the idea of proving equivalence between programs is not new. It is a rather old challenge in the theorem-proving community. A lot of attention has been given to this problem in the ACL2 community (see, e.g., [11, 12, 13]). These works are mostly concerned with program equivalence as a case study for using proof techniques that are generic (i.e., not specific for proving equivalence). It seems that no one has targeted programs with similar call graphs and large parts that are syntactically equivalent, which is the target of regression verification. Another related line of work is that of *relational verification* [14, 15], which suggests a logical framework for expressing relations between programs and corresponding deductive rules for proving them, which are targeted at popular types of transformations, e.g., loop unrolling.

Attempts to build fully automatic proof engines for equivalence of industrial programs concentrated so far on very restricted cases. Specifically, they all focused on programs without dynamic memory allocation and with bounded loops (or loop-free). Feng and Hu, for example, considered the problem of proving equivalence of embedded code [16]. The main technique used in this line of work is to prove the equivalence of small segments of the code after unrolling loops some predefined number of times. Arons et al. [17] developed a tool at Intel for proving the equivalence of two versions of microcode, with the goal of proving backwards compatibility, but the programs were assumed to be loop-free.

Another relevant line of research is concerned with *translation validation* [18, 19], the process of proving equivalence between a source and a target of a compiler or a code generator. The fact that the translation is mechanical allows the verification methodology to rely on various patterns and restrictions on the generated code. For example, translation validation for synchronous languages [19, 20] relies on the fact that the target C code has exactly one loop (corresponding to a *step* in the synchronous program) and hence the proof is conducted by induction over an expression which is derived from loop-free code. A recent example of translation validation, from the synchronous language SDL to C, is by Haroud and Biere [21]. It is based on a variation of Floyd’s method [22] for proving equivalence: it declares cutpoints in both programs (as in the original Floyd’s method, there should be at least one cutpoint in each loop), maps them between the two programs, and proves that two related cutpoints are equivalent with respect to the ‘observable’ variables if they are equivalent in the preceding pair of cutpoints. This method works in the boundary of a single function on each side and does not support general programs (e.g., recursive programs).

[§]Based on an early version of this article, the algorithm suggested here has recently been implemented in SymDiff.

2. PROVING PARTIAL EQUIVALENCE

We begin by recalling Hoare’s rule for recursive invocation. Let $\{p\} \text{ call } f \{q\}$ be a Hoare triple where f is a recursive function. The inference rule suggested by Hoare [23] for proving this triple is:

$$\frac{\{p\} \text{ call } f \{q\} \vdash_H \{p\} f \text{ body } \{q\}}{\{p\} \text{ call } f \{q\}} \text{ (REC) ,} \quad (1)$$

where “ f body” is the body of f . In words, the only effect of the recursive call on the proof is that it is assumed that it maintains the (p, q) relation. This unintuitive rule was described by Hoare [23] as follows: *The solution... is simple and dramatic: to permit the use of the desired conclusion as a hypothesis in the proof of the body itself.* The correctness of rule (REC) is proved by induction, where the base case corresponds to the base(s) of the recursion, namely the nonrecursive run(s) through the procedure.

2.1. Rule (PROC-P-EQ_s)

The rule for partial equivalence between functions f and f' has the same flavor as (REC). It applies to the special case of f, f' being recursive functions without calls to other functions. It is called (PROC-P-EQ_s), for ‘Procedures Partial Equivalence’, where the subscript s indicates that it applies only to the special case.

$$\frac{p\text{-equiv}(\text{call } f, \text{call } f') \vdash p\text{-equiv}(f \text{ body}, f' \text{ body})}{p\text{-equiv}(\text{call } f, \text{call } f')} \text{ (PROC-P-EQ}_s\text{) .} \quad (2)$$

Informally, this means that if assuming that the recursive calls are partially equivalent enables us to prove this condition over the bodies of f and f' , then f and f' are partially equivalent. It was shown in a previous publication [7] that this rule is sound. Although the soundness proof refers to an artificial abstract language, it has most of the features of an imperative language such as C. In Sect. 3.7 this point will be elaborated further.

A convenient method for checking the premise of rule (PROC-P-EQ_s) is to replace the recursive call with an *uninterpreted function* because by definition instances of the same uninterpreted function are partially equivalent (this is guaranteed by the congruence axiom that defines such functions [24]). After performing this replacement the calling function is said to be *isolated*. Denote by f^{UF} the isolated version of a function f . Rule (PROC-P-EQ_s) can be reformulated accordingly:

$$\frac{p\text{-equiv}(\text{call } f^{UF}, \text{call } f'^{UF})}{p\text{-equiv}(\text{call } f, \text{call } f')} . \quad (3)$$

Since every loop can be extracted to a recursive function, and the recursive functions are replaced with uninterpreted functions, the premise of this rule is decidable for a language with finite domains such as C. The following example demonstrates the use of this rule.

Example 1

Consider the two functions in Fig. 1. Let U be the uninterpreted function such that calls to U replace the recursive calls to `gcd1` and `gcd2`. Figure 2 presents the isolated functions.

To prove the partial equivalence of the two functions, it is necessary to first translate them to formulas expressing their respective transition relations. A convenient way to do so is to use Static Single Assignment (SSA) [24]. Briefly, this means that in each assignment of the form $x = \text{exp}$; the left-hand side variable x is replaced with a new variable, say x_1 . Any reference to x after this line and before x is assigned again is replaced with the new variable x_1 (this is done in a context of a program without unbounded loops). In addition, assignments are guarded according to the control flow. After this transformation, the statements are conjoined: the resulting equation represents the computations of the original program.

The SSA form of isolated `gcd1`, denoted T_{gcd1} , is

```

gcd1(int a, int b)          gcd2(int x, int y)
{ int g;                    { int z;
  if (!b) g = a;            z = x;
  else {                    if (y > 0)
    a = a%b;                z = gcd2(y, z%y);
    g = gcd1(b, a); }      return z;
  return g;                 }
}

```

Figure 1. Two functions to calculate the GCD of two nonnegative integers.

```

gcd1(int a, int b)          gcd2(int x, int y)
{ int g;                    { int z;
  if (!b) g = a;            z = x;
  else {                    if (y > 0)
    a = a%b;                z = U(y, z%y);
    g = U(b, a); }          return z;
  return g;                 }
}

```

Figure 2. After isolation of the functions, i.e., replacing their function calls with calls to the uninterpreted function U .

$$\left(\begin{array}{l} a_0 = a \\ b_0 = b \\ b_0 = 0 \rightarrow g_0 = a_0 \\ (b_0 \neq 0 \rightarrow a_1 = (a_0 \% b_0)) \wedge (b_0 = 0 \rightarrow a_1 = a_0) \\ (b_0 \neq 0 \rightarrow g_1 = U(b_0, a_1)) \wedge (b_0 = 0 \rightarrow g_1 = g_0) \\ g = g_1 \end{array} \wedge \right). \quad (4)$$

The SSA form of isolated gcd2 , denoted T_{gcd_2} , is

$$\left(\begin{array}{l} x_0 = x \\ y_0 = y \\ z_0 = x_0 \\ y_0 > 0 \rightarrow z_1 = U(y_0, (z_0 \% y_0)) \\ y_0 \leq 0 \rightarrow z_1 = z_0 \\ z = z_1 \end{array} \wedge \right). \quad (5)$$

The premise of rule (PROC-P-EQ_s) requires proving the validity of the following formula over nonnegative integers:

$$(a = x \wedge b = y \wedge T_{\text{gcd}_1} \wedge T_{\text{gcd}_2}) \rightarrow g = z. \quad (6)$$

Many theorem provers can prove such formulas fully automatically, and hence establish the partial equivalence of gcd1 and gcd2 . Note that a real implementation of the GCD function should include a check that the two input numbers are non-negative, and return some error code otherwise. Under this modification, the two functions can be proved to be partially equivalent without the necessity to assume that the input is non-negative. \square

2.2. Extensions to (PROC-P-EQ_s)

Now suppose that the two compared functions f, f' call other functions f_c, f'_c , respectively. Rule (PROC-P-EQ_s) can still be used if one of the following holds:

1. If f_c, f'_c were already proven to be equivalent then they can be replaced with uninterpreted functions. Such a replacement imposes an overapproximating abstraction. The soundness of the rule is maintained.
2. Otherwise, if f_c, f'_c and their descendants are not recursive then they can be inlined in their callers. The premise of rule (PROC-P-EQ_s) is then checked as before.
3. Otherwise, if some of the descendants of f_c, f'_c are recursive but were proven partially equivalent then these descendants can be abstracted with uninterpreted functions. As in the previous case f_c, f'_c can then be inlined into their callers.

2.3. A generalization: rule (PROC-P-EQ)

(PROC-P-EQ_s) can be generalized to mutually recursive functions. In the call graphs these appear as strongly connected components (SCCs) of size larger than one. The focus here is on Maximal SCCs, or MSCCs. Nodes that are not part of any cycle in the call graph correspond to nonrecursive functions, and are called *trivial* MSCCs. Consider two nontrivial MSCCs m, m' in the two programs. Assume that the functions in m, m' do not have loops nor do they call functions outside of m and m' . Further assume that there is a bijective mapping map_f between the functions in m and m' . The generalization of (PROC-P-EQ_s) to mutually recursive functions is:

$$\frac{\forall(f, f') \in map_f. ((\forall(g, g') \in map_f. p\text{-equiv}(\text{call } g, \text{call } g')) \vdash p\text{-equiv}(f \text{ body}, f' \text{ body}))}{\forall(f, f') \in map_f. p\text{-equiv}(\text{call } f, \text{call } f')} \quad (\text{PROC-P-EQ}) . \quad (7)$$

The version of this rule with uninterpreted functions is defined next. For a function g , let $UF(g)$ be an uninterpreted function such that g and $UF(g)$ have the same prototype. The following relation is enforced: $(g, g') \in map_f \Leftrightarrow UF(g) = UF(g')$. Let

$$f^{UF} \doteq f[g \leftarrow UF(g) \mid g \text{ is called in } f] .$$

The generalization of (3) is:

$$\frac{\forall(f, f') \in map_f. p\text{-equiv}(f^{UF}, f'^{UF})}{\forall(f, f') \in map_f. p\text{-equiv}(\text{call } f, \text{call } f')} . \quad (8)$$

The reader may observe the resemblance of this rule to that of Hoare's rule for mutual recursion [25].

Example 2

Consider the two small MSCCs in the top part of Fig. 3, where $map_f = \{(f_1, f'_1), (f_2, f'_2)\}$. According to (8) the partial equivalence of (f_1, f'_1) has to be proven while replacing the calls to f_2, f'_2 with the same uninterpreted function ($U_2 = UF(f_2) = UF(f'_2)$, see middle drawing), and the partial equivalence of (f_2, f'_2) has to be proven separately while replacing the calls to f_1, f'_1 with the same uninterpreted function ($U_1 = UF(f_1) = UF(f'_1)$, see bottom drawing). □

2.4. Extensions and relaxations of (PROC-P-EQ)

Now suppose that functions in m and m' do call functions outside of m and m' , respectively. All the extensions listed in Sect. 2.2 for the case of simple recursive functions apply here (e.g., nonrecursive functions can be inlined). Further, suppose that there is no bijective mapping between the functions in m and m' , or that not all pairs in map_f are partially equivalent (or can be proven to be so). It is now shown that it still may be possible to prove the partial equivalence of *some* of the functions. The idea is to inline some of the functions in m and m' , as long as cycles are not created. Formally, let $S \subseteq \{\langle f, f' \rangle \mid \langle f, f' \rangle \in map_f, f \in m, f' \in m'\}$ be a set of function pairs that satisfies:

- Every cycle in m contains a node f such that $\exists f'. \langle f, f' \rangle \in S$, and
- Every cycle in m' contains a node f' such that $\exists f. \langle f, f' \rangle \in S$.




Figure 3. Proving partial equivalence of mutual recursive functions.

The rule for proving the partial equivalence of the pairs in S is then:

$$\frac{\forall (f, f') \in S. p\text{-equiv}(f^{UF}, f'^{UF})}{\forall (f, f') \in S. p\text{-equiv}(\text{call } f, \text{call } f')}, \quad (9)$$

but here the definition of f^{UF} and f'^{UF} is slightly different than the definition in Sect. 2.1. Let m_S denote all the functions in m that participate in S , i.e., $m_S \doteq \{g \mid g \in m \wedge \exists g'. (g, g') \in S\}$. m'_S is defined similarly with respect to m' . f^{UF} is now created by replacing each call to a function $h \in m_S$ by a call to $UF(h)$, and inlining the rest. f'^{UF} is created similarly. Hence, the difference from the original definition of f^{UF} and f'^{UF} is that now the functions outside of $\{m_S \cup m'_S\}$ are inlined.

The rule indicates that the requirement of bijectiveness of map_f can be relaxed: it can be any 1-1 partial mapping that contains a set S as defined above. Note that S replaces map_f in both the premise and the consequent of (PROC-P-EQ). Hence the weaker premise has a weaker consequent. Ideally S should be as close as possible to the original definition, namely bijective, because then more pairs of functions are proven to be partially equivalent.

In the next section an algorithm is described, which attempts to prove partial equivalence of general programs by traversing the call graphs bottom-up and replacing functions with their uninterpreted versions when possible, based on these generalizations.

3. A DECOMPOSITION ALGORITHM

3.1. Preliminaries: CBMC

RVT is geared towards C programs, and uses CBMC [26] as the underlying decision procedure for checking the premise of rule (PROC-P-EQ) and its extensions as described in the previous section. CBMC, developed by D. Kroening, is a bounded model checker for C programs that supports almost all of the features of ANSI-C. It requires from the user to define a bound k on the number of iterations that each loop in a given ANSI-C program is taken, and a similar bound on the depth of each recursion. It also provides functions that return nondeterministic values (e.g., `nondet_int()` and `nondet_float()`), which are used for modeling the possible inputs. This enables CBMC to symbolically characterize the full set of possible executions restricted by the user-defined bounds, by a decidable formula f (by default f is propositional, but CBMC can generate the verification condition in other decidable theories). The existence of a solution to $f \wedge \neg a$, where a is a user defined assertion, implies the existence of a path in the program that violates a . Otherwise, CBMC is said to have established the k -correctness of the checked assertions. CBMC is used here in a very restricted way, however: recall that the premise of rule (PROC-P-EQ) is over nonrecursive functions without loops, hence in this case $k = 1$.

RVT generates small loop-free and recursion-free C programs – each corresponds to a pair of functions that it attempts to prove equal – which it sends to CBMC for decision.

3.2. Preprocessing and mapping

Before this iterative process begins, RVT makes two preliminary steps.

Loops All loops in f and f' are replaced with recursive functions. This process is described in the first author's thesis [9].

Mapping A (possibly partial) mapping map is built by pairing functions and global variables between the two compared programs. Mapping is done recursively, in a manner reminiscent of computing congruence closure. The algorithm works on the parse trees of both programs and pairs nodes, where a node can be either a variable, a function, or a type. Initially it maps global variables with the same name and type. It then maps functions with the same name, return type, prototype, and such that their lists of global variables that they read and write to are mapped pairwise. Then, within mapped functions that are also syntactically equivalent up to variable names, it attempts to map elements that appear in isomorphic locations. If these elements were already mapped it just checks that the mapping according to this function agrees with the previous one, and otherwise it issues a warning. This process is repeated until no new mapping is discovered.

Note that wrong mapping does not affect soundness: it is used for generating the verification conditions, and hence wrong mapping can only fail a proof.

Denote by map_f the partial mapping of functions resulting from this process. For simplicity of the presentation assume from hereon that the global variables accessed by a function are added at the end of its list of parameters, in a consistent order. This assumption simplifies the description of the algorithm later on.

The input for the main algorithm is thus two recursive programs without loops, and a mapping between the functions map_f .

3.3. A bottom-up decomposition algorithm

The equivalence check in RVT is presented in Algorithm 1. It is based on traversing bottom-up the call graphs of the two programs to be compared. In line 2 all nonrecursive functions that are not mapped are inlined. In the next line the MSCC DAGs MD_1 and MD_2 are built from the call graphs of the input programs. An MSCC DAG corresponding to a program is simply the call graph of the program after collapsing its MSCCs into single nodes. In line 4 the algorithm attempts to build a bijective mapping map_m between the nodes of MD_1 and MD_2 , which is consistent with map_f . In other words, if $\langle m_1, m_2 \rangle \in map_m$, f is a function in m_1 and $\langle f, f' \rangle \in map_f$, then f' is a function in m_2 (and vice-versa). If such a mapping is impossible, the algorithm aborts. In practice one may run the algorithm bottom-up until reaching nonmapped MSCCs, but this option is omitted here in order to keep the description simple.

In line 5 the bottom-up traversal begins. The algorithm searches for the next unmarked pair of MSCCs such that its children pairs are already marked. If the selected pair $\langle m_1, m_2 \rangle$ is trivial, then in line 9 the equivalence of the two functions in m_1, m_2 is checked. The function CHECK is described in Alg. 2. Otherwise, namely $\langle m_1, m_2 \rangle$ is not trivial, the algorithm proceeds in line 11 by choosing nondeterministically a subset S of paired functions from m_1, m_2 that intersect all cycles in m_1 and m_2 (in graph-theoretic terms, the functions in S constitute a *feedback vertex set* of both m_1 and m_2). The algorithm can be determinized by, e.g., attempting all such sets.[¶] A good strategy, as implied by the discussion in the end of Sect. 2, is to give priority to larger sets, since the larger the set is, the more functions are proven to be partially equivalent. Further, larger sets imply less functions to inline, and hence the burden on the decision procedure is expected to be smaller. RVT solves this optimization problem by reducing it to a Pseudo-Boolean formula and invoking an off-the-shelf PBS solver. If one of the pairs in S cannot be proven to be equivalent, the algorithm must

[¶]Although there can be an exponential number of them in the size of the MSCC, observe that large MSCCs in real programs are rare.

abort: neither $\langle m_1, m_2 \rangle$ nor the SCCs above it can be proven equivalent by the algorithm. Otherwise in line 14 all the functions that are paired in S are marked as “Equivalent”. Finally, $\langle m_1, m_2 \rangle$ is marked as “Covered”, and the algorithm continues to the next pair.

Algorithm 1 A bottom-up decomposition algorithm for proving the partial equivalence of pairs of functions.

```

1: function PROVE(Programs  $P, P'$ , map between functions  $map_f$ )
2:   Inline nonrecursive nonmapped functions;
3:   Generate MSCC DAGs  $MD_1, MD_2$  from the call graphs of  $P, P'$ ;
4:   If possible, generate a bijective map  $map_m$  between nontrivial nodes in  $MD_1$  and  $MD_2$ 
      that is consistent with  $map_f$  (it is desirable but not necessary to add pairs of trivial nodes
      to  $map_m$ ). Otherwise abort.
5:   while  $\exists \langle m_1, m_2 \rangle \in map_m$  that is uncovered and its children are “Covered” do
6:     Choose such a pair  $\langle m_1, m_2 \rangle$ ;
7:     if  $m_1, m_2$  are trivial then
8:       Let  $f_1, f_2$  be the functions in  $m_1, m_2$ , respectively;
9:       if CHECK( $f_1, f_2$ ) then mark  $f_1, f_2$  as “Equivalent”;
10:    else
11:      Select nondeterministically a set of function pairs  $S \subseteq \{\langle f, f' \rangle \mid \langle f, f' \rangle \in map_f,$ 
         $f \in m_1, f' \in m_2\}$  that intersect all cycles in  $m_1$  and  $m_2$ ;
12:      for all  $\langle f, f' \rangle \in S$  do
13:        if  $\neg$ CHECKr( $f, f', S$ ) then abort;
14:        for all  $\langle f, f' \rangle \in S$  do mark  $f, f'$  as “Equivalent”;
15:      Mark  $\langle m_1, m_2 \rangle$  as “Covered”.

```

Now consider CHECK, which appears in Alg. 2. This function begins by checking whether the input functions happen to be syntactically equivalent and their children are also marked equivalent. If yes, it returns true. Otherwise it sends CBMC a nonrecursive, loop-free C program, which is called here a *check-block*. Following is a description of this program.

Let D_f denote the maximal connected subDAG rooted at f that contains only functions that are unpaired or *not* marked “Equivalent”, but excluding f itself. $D_{f'}$ is defined similarly with respect to f' . The program *check-block* (f, f') consists of the following elements:

1. The functions f, f' and all functions in $D_f, D_{f'}$, such that
 - Name collisions in global identifiers of the two programs are solved by renaming;
 - All calls to f, f' are replaced with calls to $UF(f)$ ($=UF(f')$), respectively;
 - For all $\langle h_1, h_2 \rangle \in map_f$ such that $h_1, h_2 \notin D_f \cup D_{f'}$, calls to h_1, h_2 are replaced with calls to $UF(h_1)$ ($=UF(h_2)$). (Observe that the pair $\langle h_1, h_2 \rangle$ must be marked “Equivalent”).
2. The `main()` function, which consists of:
 - Assignment of nondeterministic but equal values to inputs of f and f' ;
 - Calls to f, f' ; and
 - Assertion that the outputs of f and f' are equal.

Following are several notes on the definition of *check-block* (f, f'):

- *check-block* is guaranteed to be nonrecursive. This is because when PROVE fails to prove the equivalence of MSCCs it aborts in line 13, and hence recursive functions that are not proven equivalent will never be part of future *check-blocks*.
- The code of each nonrecursive pair $\langle f, f' \rangle \in map_f$ that could not be proven equivalent is included when checking the equivalence of their parents, and possibly more ancestors, until reaching a provably equivalent pair or reaching the roots. This process is called here *logical inlining*, since it is equivalent to inlining but is more faithful to the program’s original

structure. This enables RVT to prove equivalence in case, for example, that some code was moved from the parent to the child, but together they still perform the original computation.

- The code of a pair $\langle f, f' \rangle \in \text{map}_f$ that is proven to be equivalent does not participate in any subsequent check-block. It is replaced with uninterpreted functions in all subsequent checks, or disappears altogether if some ancestor pair is also marked “Equivalent” in each of its paths to the roots of the related subprograms.
- The replacement of recursive calls of paired functions with uninterpreted functions corresponds to isolation (see Sect. 2). Recall that proving equivalence of mapped isolated functions also proves their partial equivalence by rule (PROC-P-EQ).

Algorithm 2 A function called by PROVE for checking the equivalence of two input nonrecursive functions. `check-block` is a C program defined in the main text.

```

1: function CHECK(function  $f$ , function  $f'$ )
2:   if  $f$  and  $f'$  are syntactically equivalent and all their children are marked “Equivalent” then
3:     return true;
4:   return CBMC (check-block ( $f, f'$ ));

```

Algorithm 3 A function called by PROVE for checking the equivalence of two input functions that are part of MSCCs. `check-blockr` is a C program defined in the main text.

```

1: function CHECKr(function  $f$ , function  $f'$ , set of pairs  $S$ )
2:   if  $f$  and  $f'$  are syntactically equivalent and all their children are either marked “Equivalent”
   or in  $S$  then
3:     return true;
4:   return CBMC (check-blockr ( $f, f', S$ ));

```

The function `CHECKr`, described in Alg. 3, is similar to `CHECK`, and can be seen as its generalization to the case that $S \neq \emptyset$. It also begins by checking for syntactical equivalence, but permits children of the checked functions to be in S rather than being marked equivalent. If this check fails, it calls `CBMC` with a program that is identical to `check-block` (f, f') as defined above, except the following difference in the definition of D_f and $D_{f'}$. Recall that D_f and $D_{f'}$ include all nodes in the subDAG under f and f' respectively that are unpaired or not marked “Equivalent”. But here, since f and f' are recursive (they are part of MSCCs), their descendants form general graphs rather than DAGs. D_f and $D_{f'}$ are now redefined so they do *not* include functions from S , which forces them to be nonrecursive. More formally, Let D_f^r be the set of functions in the maximal connected graph descending from f which does *not* include

- f itself,
- functions in S , and
- functions that are marked “Equivalent”.

$D_{f'}^r$ is defined similarly with respect to f' . Hence D_f^r and $D_{f'}^r$ are nonrecursive. Replace $D_f, D_{f'}$ with $D_f^r, D_{f'}^r$ in the definition of `check-block` to get `check-blockr`.

3.4. Examples

In this section Alg. 1 is demonstrated with two examples. The first focuses on programs with simple recursion only, and the second on the more general case.

Example 3

Consider the call graphs in Fig. 4. Assume that for $i = 1, \dots, 6$ $\langle f_i, f'_i \rangle \in \text{map}_f$, and that the functions marked by gray nodes in Fig. 4 are syntactically equivalent to their counterparts.

The execution of Algorithm 1 is described next, step by step. Initially, in line 2, f'_7 is inlined into f'_2 . The iterations that follow are listed below: □




Figure 4. Two call graphs for Example 3. A node is gray if it is syntactically equivalent to its counterpart.

It.	Pair	Description	Res.
1.	$\langle f_3, f'_3 \rangle$	Since f_3, f'_3 are nonrecursive CHECK is called in line 9, which returns TRUE based on a syntactic check.	✓
2.	$\langle f_4, f'_4 \rangle$	Same as with $\langle f_3, f'_3 \rangle$.	✓
3.	$\langle f_6, f'_6 \rangle$	Handled by CHECK, which calls CBMC. Assume that CHECK returns FALSE.	×
4.	$\langle f_2, f'_2 \rangle$	Handled by CHECK, which calls CBMC. check-block contains f_2, f'_2 (recall that f'_7 is inlined into f'_2), and the calls to f_3, f'_3, f_4, f'_4 are replaced by calls to uninterpreted functions. Assume that this check fails, and hence CHECK returns FALSE.	×
5.	$\langle f_5, f'_5 \rangle$	Handled by CHECK ^r . The only choice for S is $\{\langle f_5, f'_5 \rangle\}$. In this case $D_{f_5}^r = \{f_6\}$ and $D_{f'_5}^r = \{f'_6\}$, hence check-block ^r contains f_5, f'_5, f_6, f'_6 , and the recursive calls are replaced with uninterpreted functions. Assume that the check succeeds.	✓
6.	$\langle f_1, f'_1 \rangle$	Handled by CHECK which calls CBMC. check-block contains f_1, f'_1, f_2, f'_2 (f'_7 is inlined into f'_2), and the calls to $f_3, f'_3, f_4, f'_4, f_5, f'_5$ are replaced by calls to uninterpreted functions.	✓




Figure 5. Two call graphs for Example 4

An example of programs with mutual recursion follows.

Example 4

Consider the call graphs that are presented in Fig. 5. Assume that for $i = 1, \dots, 6$ $\langle f_i, f'_i \rangle \in \text{map}_f$. As in the previous example syntactically equivalent functions are marked by gray nodes. The nodes of the MSCC DAGs that are generated in line 3 are (listed bottom-up, left-to-right): $MD_1 = \{\{f_3\}, \{f_4\}, \{f_6\}, \{f_2, f_5\}, \{f_1\}\}$ and $MD_2 = \{\{f'_3\}, \{f'_4\}, \{f'_6\}, \{f'_2, f'_5, f'_7\}, \{f'_1\}\}$. The MSCC mapping map_m in line 4 is naturally derived from map_f . The first three iterations, over $\langle f_3, f'_3 \rangle$, $\langle f_4, f'_4 \rangle$ and $\langle f_6, f'_6 \rangle$, are the same as in the previous example. Next comes the nontrivial SCCs. Assume that in line 11 $S = \{\langle f_2, f'_2 \rangle, \langle f_5, f'_5 \rangle\}$ is chosen. This is more than strictly necessary for breaking the cycles, but recall that the aim is to prove partial equivalence for as many pairs as possible. The iterations from there on are listed below. Denote by \checkmark^c that the equivalence check succeeded, but the functions are not yet marked as equivalent:

It.	Pair	Description	Res.
4.	$\langle f_2, f'_2 \rangle$	Handled by CHECK ^r , which calls CBMC. check-block ^r contains f_2, f'_2, f'_7 . Calls to f_3, f'_3, f_4, f'_4 and f_5, f'_5 are replaced with calls to uninterpreted functions. Assume that this check succeeds.	\checkmark^c
5.	$\langle f_5, f'_5 \rangle$	Handled by CHECK ^r , which calls CBMC. check-block ^r contains f_5, f'_5, f_6, f'_6 , and the calls to f_2, f'_2 are replaced with calls to uninterpreted functions. Assume that the check succeeds. (In line 14 $\langle f_2, f'_2 \rangle$ and $\langle f_5, f'_5 \rangle$ are marked “Equivalent”).	\checkmark^c
6.	$\langle f_1, f'_1 \rangle$	Handled by CHECK. The syntactic check returns TRUE. (All MSCCs are marked “Covered” and the algorithm terminates).	\checkmark

What would happen had $S = \{\langle f_2, f'_2 \rangle\}$ was chosen in line 11? In that case check-block^r would contain $f_2, f'_2, f'_7, f_5, f'_5, f_6, f'_6$, and calls to f_3, f'_3, f_4, f'_4 and f_2, f'_2 would be replaced with calls to uninterpreted functions. If this check would have succeeded, then only $\langle f_2, f'_2 \rangle$ would be marked as equivalent. Not only that this would not prove the equivalence of $\langle f_5, f'_5 \rangle$, but it also would complicate the next step: when checking the pair $\langle f_1, f'_1 \rangle, f_5, f'_5, f_6, f'_6$ would have to be included, as they were not marked equivalent. Calls to $\langle f_2, f'_2 \rangle$ would be replaced with calls to uninterpreted functions. This example motivates the suggested determinization: try large S first. \square

3.5. Controlling the abstraction level

The definition of check-block^r entails that every call to a function in S is replaced with a call to an uninterpreted function. This is not always necessary. Replacing functions with uninterpreted functions entail less computational effort, but also loses precision, and hence may fail proofs. Some control over the abstraction level is possible even for a given set S , because the only thing that needs to be guaranteed is that check-block^r is nonrecursive, and that every pair of functions that are assumed to be partially equivalent is eventually proven to be so. Consider once again Example 4, where the set $S = \{\langle f_2, f'_2 \rangle, \langle f_5, f'_5 \rangle\}$ was chosen. When checking $\langle f_5, f'_5 \rangle$, the calls to f_2, f'_2 were replaced with uninterpreted functions because $\langle f_2, f'_2 \rangle \in S$. Instead these functions and f'_7 can be included in check-block^r, and the calls to f_5, f'_5 can be replaced with calls to uninterpreted functions.

It is obvious that the actual functions matter for choosing the most refined abstraction. Without this information a reasonable heuristic is to maximize the number of calls to interpreted functions, while still breaking all cycles.

3.6. Correctness of PROVE

Theorem 1

(Correctness of PROVE) Functions that are marked “Equivalent” by Alg. 1 are partially equivalent to their counterpart in map_f .

Proof

A proof sketch is given, based on the soundness of (PROC-P-EQ) and its generalization in Eq. (9), and on the soundness of CBMC itself. The program sent to CBMC clearly checks for partial equivalence; hence the correctness argument for the case CBMC is invoked must show that the abstraction induced by replacing functions with uninterpreted functions is a conservative one.

Consider the current covering (the result of executing line 15) of MD_1 nodes right after line 15. Let d denote the current covering ‘depth’, i.e., the largest distance of a covered node from any leaf node in MD_1 . The proof is by induction on d .

Base: When $d = 1$, m_1 is a leaf. There are two cases to consider, corresponding to the lines in which the marking of functions with “Equivalent” is done:

- m_1 and m_2 are trivial nodes, and the single function that m_1 contains, denoted f , is marked “Equivalent”. It must have been marked by CHECK in line 9. CHECK returns TRUE if f and its counterpart are either syntactically equivalent or proven to be equivalent by CBMC. In the latter, check-block does not include calls to uninterpreted functions.

- m_1 (or m_2) is not trivial, and a subset of functions corresponding to S of line 11 is marked “Equivalent”. The condition for the marking in line 14 corresponds exactly to the premise of (9).

Step: Now assume that the theorem is correct for d , and consider $d + 1$. For a function f marked “Equivalent” in level $d + 1$, such that $\langle f, f' \rangle \in \text{map}_f$, it will be proven that f and f' are indeed partially equivalent. The proof relies on the induction hypothesis, which implies that descendant functions of f and f' that are marked “Equivalent” are indeed partially equivalent to their counterparts in map_f . As in the base case, the two lines in which the marking of functions with “Equivalent” is done, are considered separately:

- f and f' constitute trivial MSCC nodes m_1 and m_2 , respectively. f must have been marked by CHECK in line 9. CHECK returns TRUE in two cases:
 - f and f' are syntactically equivalent and their children are marked “Equivalent”. Hence f is indeed partially equivalent to f' .
 - Otherwise, CBMC must have returned TRUE. check-block contains, in addition to f and f' , either inlining of nonrecursive functions or calls to uninterpreted functions. The former clearly preserves correctness. The latter is a conservative abstraction of the original functions, which, recall, are partially equivalent by the induction hypothesis. Hence f and f' are partially equivalent.
- m_1 (or m_2) is not trivial, and a subset of functions corresponding to S is marked “Equivalent” in line 14. Let $\langle f, f' \rangle \in \text{map}_f$ be any of the pairs in S . CHECK^r (on line 3) returns TRUE in two cases:
 - f and f' are syntactically equivalent and their children are either marked “Equivalent” or in S . The former children are partially equivalent by the induction hypothesis. As for the latter children, f can only be marked “Equivalent” if all the pairs in S passed CHECK^r, which by (9) means that they are indeed partially equivalent. Hence in both cases, the fact that f and f' are marked “Equivalent” implies that they indeed are.
 - Otherwise, CBMC must have returned TRUE. check-block^r possibly contains, in addition to f and f' :
 - * Inlining of nonrecursive functions (outside of m_1 and m_2).
 - * Inlining of functions in m_1, m_2 that are not part of S , where calls to S functions are replaced with uninterpreted functions. The argument made above for syntactic checks in which S functions are replaced with uninterpreted functions apply here as well: f can only be marked “Equivalent” if all the pairs in S passed CHECK^r, which by (9) means that they are indeed partially equivalent. Hence the fact that f and f' are marked as partially equivalent implies that they indeed are.
 - * Calls to uninterpreted functions that abstract functions outside of m_1 and m_2 . By the induction hypothesis, these are indeed partially equivalent, and hence the calls to uninterpreted functions is a conservative abstraction.

□

3.7. Dynamic data structures

RVT works on C (reference ANSI C99) programs, although not all features are supported. A major issue in applying rule (PROC-P-EQ) to C programs is that of dynamic data structures. Recall that deciding formulas with uninterpreted functions requires the comparison pair-wise of the arguments with which such functions are called, and a similar comparison of their outputs. If some of these arguments are pointers, such a comparison is meaningless. In this section RVT’s method of treating pointer arguments of functions and dynamic data structures is briefly described.

Whereas in nonpointer variables the comparison is between values, in the case of pointer variables the comparison should be between the data structures that they point to. A dynamic data structure can be represented as a graph – called here a *pointer-element graph* – that its vertices are fields and

its edges connect pointer fields to their objects. A simplifying assumption is made that all dynamic structures that are passed to a function through pointer arguments or globals are in the form of trees, i.e., aliasing within dynamic structures and between function arguments is not allowed. Equality of structures is defined as follows:

Definition 3

(Iso-equal structures) Two structures are *iso-equal* if their pointer-element graphs are isomorphic and the values at structs related by the isomorphism are equal.

Let p_1, p_2 be paired pointer variables that are arguments to the functions that are compared. RVT generates two iso-equal tree-like data structures with a bounded depth (see below) and with nondeterministic values (including possible null values for pointers). It then makes p_1 and p_2 point to these trees. This guarantees that the input structure is arbitrary but equivalent up to a bound, and under the assumption that on both sides it is a tree. A similar strategy is activated when comparing p_1 and p_2 that point to an output of the compared functions – the output structures are compared up to the same bound assuming they are trees.

What should be the bound on this tree? Recall that the code of the related subprograms that are checked (the check-block) does not contain loops or recursion, and hence there is a bound on the maximal depth of the items this code can access in any dynamic data structure that is passed to the roots of the related subprograms. It is possible, then, to compute this bound, or at least overestimate it, by syntactic analysis. For example, searching for code that progresses on the structure such as $n = n \rightarrow next$ for a pointer n . However, such a mechanism is not implemented yet in RVT and it relies instead on a user-defined bound.

4. EXPERIENCE WITH RVT AND MS-SYMDIFF

4.1. A case study with RVT

Consider the two mutually-recursive programs in Fig. 6, and the output of RVT that appears in Fig. 7 (the legend is explained in the caption). RVT's output is a graphical representation of the two call graphs, annotated with the results and additional debugging information. The two functions `Loop_main_for_1()` and `Loop_main_for_2()` that appear in Fig. 7 correspond to the two `for` loops in the `main()` function that were extracted by RVT to separate functions.

There are several things to note about this example. First, observe that the two SCCs ($\{F, M\}$, $\{F, M, G\}$) are of different sizes and that $G()$ has no counterpart. Nevertheless RVT is able to prove the equivalence of $F()$ and $M()$, because they break all cycles in the SCCs, as was explained in Sec. 3.3. Second, note that the output of the first loop function (`Loop_main_for_1()`) is both `out` and the return value imposed by the `return` statement. Although the latter is different in the two sides, since the `if` condition is never satisfied, this function is still equivalent on both sides. Third, observe that RVT proved that `Loop_main_for_2()` is equivalent, although the descendant `val()` is different. It demonstrates how inlining non-recursive descendants can facilitate the equivalence proof of their callers. For the same reason RVT was able to prove the equivalence of $M()$. Finally, observe that `main()` was declared equivalent by RVT without calling CBMC, because it is syntactically equivalent and its descendants were proven equivalent. Overall RVT solves this case fully automatically in nine seconds, which include four calls to CBMC.

4.2. Other experiments with RVT

RVT was tested on several synthetic and limited-size industrial programs and attempted to prove equivalent different versions of these programs:

Random programs A random program generator was used to create several dozen recursive programs of different sizes. The user specifies the probability to generate each type of variable, block, or operator. Variables can be global, local or formal arguments of functions. Types can be

```

int M(int n);

int val(int x) {
  return x;
}

int F(int n) {
  if ( n == 0 ) return 1;
  return n - M(F(n-1));
}

int M(int n) {
  if ( n == 0 ) return 0;
  return n - F(M(n-1));
}

int main() {
  int i, out;
  for(i = 0; i < 20; i++) {
    out = F(i);
    if (i > 20) return 3;
  }
  for(i = 0; i < 20; i++) {
    out = val(M(i));
  }
  return out;
}

int M(int n);
int G(int n);

int val (int x) {
  return x + 1;
}

int F(int n){
  if ( n == 0 ) return (n == 0);
  return n - M(F(n-1));
}

int M(int n) {
  if ( n == 0 ) return 0;
  return val(n) - 1 - G(M(n-1));
}

int G(int n) { return F(n); }

int main(){
  int i, out;
  for(i = 0; i < 20; i++) {
    out = F(i);
    if (i > 25) return 12;
  }
  for(i = 0; i < 20; i++) {
    out = val(M(i)) - 1;
  }
  return out;
}

```

Figure 6. Two versions of a mutually recursive program. Although the mapped SCCs are different, RVT is able to prove the equivalence of some of the functions in it, as depicted in Fig. 7.

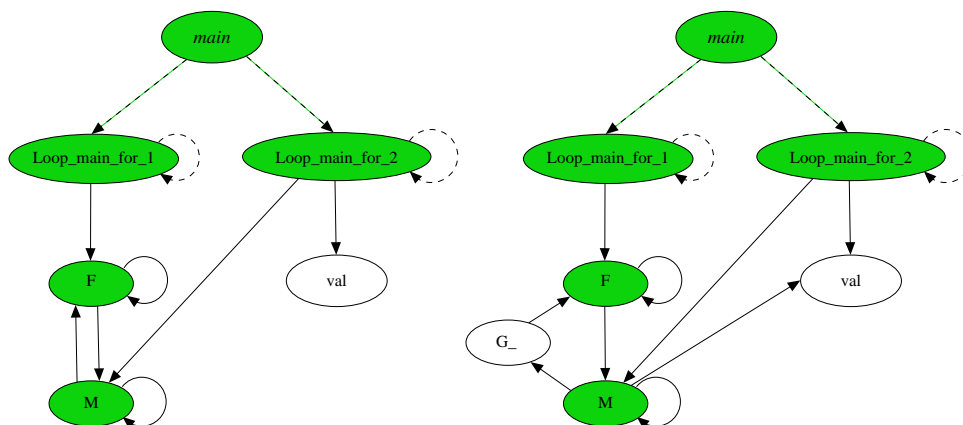


Figure 7. The output of RVT given the two programs in Fig. 6. Nodes with the same label are paired. Nodes with dark background were proven equivalent to their counterparts. Dashed edges denote a call to a loop that was extracted into a separate (recursive) function. Italicized function names (only *main* in this case) represent functions that were proven equivalent based on syntactic equivalence.

basic C types, structures or pointers to such types. Small differences between the two versions of each program are introduced in random places. This program was used to generate random yet executable C programs with up to 20 functions and thousands of lines of code. When the random versions are equivalent, RVT proves them to be partially-equivalent relatively fast, ranging from few seconds to 30 minutes. On non-equivalent versions, on the other hand, attempts to prove partial equivalence may run for many hours or run out of memory.

Industrial programs The small industrial programs that were tried are:

TCAS (Traffic Alert and Collision Avoidance System) is an aircraft conflict detection and resolution system used by all US commercial aircraft. The same 300-line fragment of this program that was used by Groce et al. [27] was used to test RVT.

MicroC/OS The core of MicroC/OS which is a low-cost priority-based preemptive real time multitasking operating system kernel for microprocessors, written mainly in C. The kernel is mainly used in embedded systems. The program is about 3000 lines long.

Matlab examples Parts of engine-fuel-injection simulation in Matlab which was generated in C from engine controller models. The tested parts contain several hundreds lines of code and use read-only arrays.

All these tests exhibit the same behavior as the random programs above. For equivalent programs, semantic-checks are very fast, proving equivalence in minutes. A case in which partially equivalent programs cannot be proven to be so due to the incompleteness of (PROC-P-EQ), were never encountered.

Recall that in the process of semantic checks, paired functions that cannot be proven equivalent are (logically) inlined. The authors' experience was that in such cases the proof becomes too hard: the decision procedure runs for hours or even fails to reach a decision at all. In some examples the bottleneck is the use of operators that burden the SAT solver, such as multiplication (*), division (/) and modulo (%) over integers. A simple solution in such cases is to *outline* these operators (i.e., take them out to a separate function). RVT then proves the equivalence of these separate functions syntactically and then replaces them with uninterpreted functions, which reduces the computation time dramatically.

4.3. Experiments with Microsoft's SymDiff

SymDiff [8] was modeled after RVT, but works at the level of the Boogie programming language [28], which means that it is language-agnostic and that its underlying proof engine is Z3 [29]. Based on an early draft of this paper, Alg. 1 was implemented in SymDiff. It replaced the original implementation that used loop and recursion unrolling, which means that it underapproximated the computations and hence could not prove equivalence in the presence of unbounded loops. Since the time that Alg. 1 was implemented in SymDiff, it was used to prove the equivalence of about 10 programs of several hundreds of lines each. In each test case manual changes were made (like deleting various lines, changing constants, outlining functions, etc), and SymDiff marked correctly the semantically different functions, typically within less than a minute.

5. SUMMARY

The introduction started by mentioning Tony Hoare's grand challenge, namely that of functional verification, and by mentioning that proving equivalence is a grand challenge in its own right, although an easier one. This work begins to explore this direction in the context of C programs, and reports on a prototype tool RVT with which the equivalence of several small industrial programs was proven. The suggested technique can be improved in several dimensions, such as strengthening rule (PROC-P-EQ) with automatically generated invariants and finding more opportunities for making the

verification conditions easier to decide. Investigating such opportunities for object-oriented code is another big challenge.

To summarize, the main contribution of this article is a method for an automatic, incremental proof, based on isolating functions from their callees and abstracting them with uninterpreted functions. This method keeps the verification conditions decidable and small relative to the size of the input programs. The initial syntactic checks and the decomposition mechanism helps meeting the goal of keeping the complexity sensitive to the changes rather than to the original size of the compared programs.

REFERENCES

1. Hoare T. The verifying compiler: A grand challenge for computing research. *J. ACM* 2003; **50**(1):63–69.
2. Dustin E. *Effective Software Testing: 50 Specific Ways to Improve Your Testing*. Addison-Wesley Professional, 2002.
3. Zeller A. Yesterday, my program worked. today, it does not. why? *ESEC / SIGSOFT FSE*, 1999; 253–267.
4. Godlin B, Strichman O. Regression verification - a practical way to verify programs. *Verified Software: theories, tools, experiments (VSTTE)*, LNCS, vol. 4171, Meyer B, Woodcock J (eds.), 2005. Conf in 2005, published in 2007.
5. Fowler M. <http://www.refactoring.com>.
6. Hoare C. An axiomatic basis for computer programming. *Comm. ACM* 1969; **12**(10):576–580.
7. Godlin B, Strichman O. Inference rules for proving the equivalence of recursive procedures. *Acta Informatica* 2008; **45**(6):403–439.
8. Kawaguchi M, Lahiri SK, Rebelo H. Conditional equivalence. *Technical Report MSR-TR-2010-119*, Microsoft Research 2010.
9. Godlin B. Regression verification: Theoretical and implementation aspects. Master's Thesis, Technion, Israel Institute of Technology 2008.
10. Godlin B, Strichman O. Regression verification. *46th Design Automation Conference (DAC)*, 2009.
11. Craciunescu S. Proving the equivalence of clp programs. *ICLP*, 2002; 287–301.
12. Manolios P, Kaufmann M. Adding a total order to acl2. *The Third International Workshop on the ACL2 Theorem Prover*, 2002.
13. Manolios P, Vroon D. Ordinal arithmetic: Algorithms and mechanization. *Journal of Automated Reasoning* 2006; To appear.
14. Benton N. Simple relational correctness proofs for static analyses and program transformations. *POPL*, 2004; 14–25.
15. Barthe G, Crespo JM, Kunz C. Relational verification using product programs. *FM*, 2011; 200–214.
16. Feng X, Hu AJ. Cutpoints for formal equivalence verification of embedded software. *EMSOFT*, 2005; 307–316.
17. Arons T, Elster E, Fix L, MadorHaim S, Mishaeli M, Shalev J, Singerman E, Tiemeyer A, Vardi MY, et al.. Formal verification of backward compatibility of microcode. *CAV05, LNCS*, vol. 3576, Springer, 2005.
18. Pnueli A, Siegel M, Singerman E. Translation validation. *TACAS'08, LNCS*, vol. 1384, Springer, 1998; 151–166.
19. Pnueli A, Siegel M, Shtrichman O. Translation validation for synchronous languages. *Proc. 25th Int. Colloq. on Automata, Languages and Programming (ICALP'98)*, LNCS, vol. 1443, Springer, 1998; 235–246.
20. Pnueli A, Siegel M, Singerman E. Translation validation. *Technical Report*, SACRES and Dept. of Comp. Sci., Weizmann Institute Apr 1997.
21. Haroud M, Biere A. SDL versus C equivalence checking. *SDL Forum*, 2005; 323–338.
22. Floyd R. Assigning meanings to programs. *Proc. Symposia in Applied Mathematics* 1967; **19**:19–32.
23. Hoare C. Procedures and parameters: an axiomatic approach. *In Proc. Sym. on semantics of algorithmic languages* 1971; (188).
24. Kroening D, Strichman O. *Decision procedures – an algorithmic point of view*. Theoretical computer science, Springer, 2008.
25. von Oheimb D. Hoare logic for mutual recursion and local variables. *Foundations of Software Technology and Theoretical Computer Science, Lecture Notes in Computer Science*, vol. 1738, Rangan C, Raman V, Ramanujam R (eds.). 1999; 168–180.
26. Kroening D, Clarke E, Yorav K. Behavioral consistency of C and Verilog programs using bounded model checking. *Proceedings of DAC 2003*, ACM Press, 2003; 368–371.
27. Groce A, Kroening D, Lerda F. Understanding counterexamples with explain. *CAV*, 2004; 453–456.
28. Barnett M, Chang BYE, DeLine R, Jacobs B, Leino KRM. Boogie: A modular reusable verifier for object-oriented programs. *FMCO*, 2005; 364–387.
29. de Moura L, Bjorner N. Z3: An efficient smt solver. *Intl. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS)*, 2008.