

Automated Discovery of Simulation Between Programs[★]

Grigory Fedyukovich¹, Arie Gurfinkel², and Natasha Sharygina¹

¹ Formal Verification Lab, University of Lugano, Switzerland,
{grigory.fedyukovich,natasha.sharygina}@usi.ch
² SEI/CMU, USA, arie@cmu.com

Abstract. Deciding equivalence between two programs (called a source and a target program) is often reduced to finding a simulation relation between them. This is computationally expensive and often requires a manual guidance. In this paper, we propose an abstraction-refinement-guided approach, called SIMABS, to automatically construct a simulation relation between the source program and an abstraction of the target program. In our approach both the abstraction and the simulation relation are discovered automatically, and deciding whether a given relation is a simulation relation is reduced to deciding validity of a $\forall\exists$ -formula. We present a novel algorithm for deciding such formulas using an SMT-solver. In addition to deciding validity, our algorithm constructs a witnessing Skolem relation. These relations enable the refinement-step of SIMABS. We have implemented SIMABS using UFO framework and Z3 SMT-solver and applied it to finding simulation relations between C programs from the Software Verification Competition. Our empirical results show that SIMABS is efficient at finding a simulation relation.

1 Introduction

There is a growing interest in the problems of regression verification and program equivalence checking [21, 22, 11, 14, 10, 20, 9]. In general, the problem is to identify (and check) the condition under which two programs (referred to as the *source* (S) and the *target* (T)) are equivalent (i.e., satisfy the same properties). These approaches prevent the wasted efforts in re-analyzing equivalent parts of the programs. For instance, while proving safety of two closely related programs, obtaining a proof of one program and adapting it to another program can be more efficient than proving each program from scratch (e.g., [10, 9]).

For example, in [9] we applied the idea of adapting proofs to analyze whether compiler optimizations preserve safety properties. While efficient, this method had a number of limitations. The most crucial one is that it required a *mapping* between variables of S and T that was either guessed automatically or provided

[★] This material is based upon work funded and supported by the Department of Defense under Contract No. FA8721-05-C-0003 with Carnegie Mellon University for the operation of the Software Engineering Institute, a federally funded research and development center. Any opinions, findings and conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the United States Department of Defense. This material has been approved for public release and unlimited distribution. DM-0001771

Report Documentation Page

Form Approved
OMB No. 0704-0188

Public reporting burden for the collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to a penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

1. REPORT DATE 18 OCT 2014		2. REPORT TYPE N/A		3. DATES COVERED	
4. TITLE AND SUBTITLE Automated Discovery of Simulation Between Programs				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Arie Gurfinkel Grigory Fedukovich /Natasha Sharygina				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Software Engineering Institute Carnegie Mellon University Pittsburgh, PA 15213				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES)				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for public release, distribution unlimited.					
13. SUPPLEMENTARY NOTES The original document contains color images.					
14. ABSTRACT					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT unclassified	b. ABSTRACT unclassified	c. THIS PAGE unclassified			

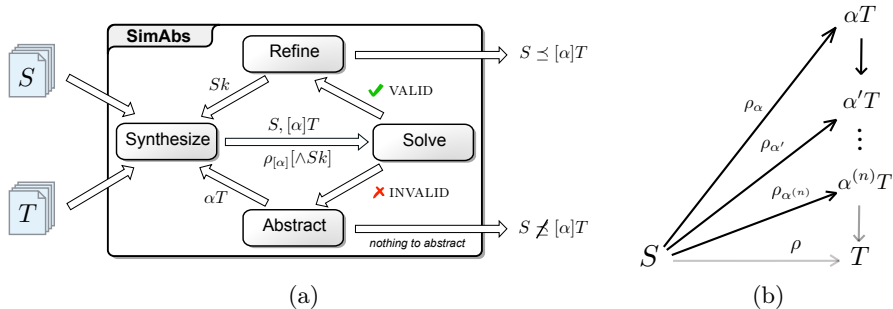


Fig. 1: (a) SIMABS and (b) its search space.

by the user. The simplest mapping that equates variables of S and T based on their names (i.e., variable a of S is mapped to a variable a of T) used in [9] is often insufficient. In practice, it sometimes results in no interesting facts lifted from S to T . As an example, consider compiler spilling that may introduce many new variable names in T that did not appear in S .

Namjoshi et al. [19, 20] show that a simulation relation is the most general mapping to transfer proofs between S and T . However, discovering a simulation relation is difficult (e.g., [20] expects the relation to be provided by the user). Moreover, their result only applies when T actually simulates S , whereas we are interested in cases where T is obtained by a small modification from S , but might not simulate S .

In this paper, we address two problems: (1) the challenge of automatically constructing a simulation relation for two arbitrary programs, and (2) if the target T does not simulate the source S , the challenge of finding a strong abstraction of the target T that simulates the source S .

Our main contribution is an iterative abstraction-refinement approach called SIMABS, illustrated in Fig. 1a. The inputs are the source and the target programs S and T , respectively. The output is an abstraction of T (possibly T itself) that simulates S and a simulation relation ρ , or a step of S that cannot be simulated by any abstraction of T (in the considered space of abstractions). SIMABS first guesses a relation ρ between S and T (**Synthesize** step). Second, it checks whether ρ is a simulation relation between S and the current abstraction of T (**Solve** step). Third, if the check succeeds, it refines the current abstraction of T and synthesizes new relation ρ (**Refine** step). Otherwise, it looks for a better abstraction α of T (**Abstract** step). The algorithm terminates when either no refinement or no abstraction is possible. The search space of the algorithm is shown in Fig. 1b. SIMABS explores the space of abstractions of T , starting with the most general abstraction (called αT in Fig. 1b) that simulates S , and iteratively refines it until no further refinement is possible.

In contrast to existing algorithms for checking whether a given relation ρ is a simulation relation (e.g., [18, 13, 21]), we reduce the problem to deciding va-

validity of a $\forall\exists$ -formula. Intuitively, the formula says “for each behavior of S there exists a corresponding behavior of $[\alpha]T$ ”. Our second contribution is a novel decision procedure, AE-VAL, for deciding validity of $\forall\exists$ -formulas over Linear Real Arithmetic (LRA). Our procedure is similar to the QE_SAT algorithm of [23]. However, in addition to deciding validity, we also extract a Skolem relation to witness the existential quantifier. This Skolem relation (called Sk in Fig. 1a) is the key to the **Refine** step of SIMABS.

We implemented SIMABS and AE-VAL on the top of the UFO framework [1, 15] and an SMT-solver Z3 [8], respectively. We have evaluated SIMABS by discovering simulation relations between programs and their optimizations (done by LLVM). Our results show that SIMABS is able to efficiently discover nontrivial simulation relations between the original and the optimized versions for most of the benchmarks.

The rest of the paper is structured as follows. After defining notation in Sect. 2, we describe how to reduce the problem of checking a simulation relation to a validity check of a specific $\forall\exists$ -formula in Sect. 3. The algorithm AE-VAL designed to solve such formulas and extract Skolem relation is presented in Sect. 4. Sect. 5 provides implementation details for the algorithm SIMABS. Sect. 6 presents an evaluation of our implementation of SIMABS. Sect. 7 compares our work with the related one, and Sect. ?? concludes the paper.

2 Background

In this section, we give the necessary definitions of a program, a transition system, and a simulation relation.

Definition 1. A program P is a tuple $\langle Var, Init, Tr \rangle$, where $Var \equiv V \cup L \cup V'$ is a set of current, next-state, and local variables, respectively; $Init$ is a formula over V that defines the initial state, and Tr is a formula over Var that denotes the transition relation.

Definition 2. A program $P = \langle Var, Init, Tr \rangle$ induces a transition system $\mathcal{T} = \langle \mathcal{S}, \mathcal{I}, \mathcal{R} \rangle$, where \mathcal{S} is a set of valuations to all variables in V (i.e., states), $\mathcal{I} = \{ \vec{s} \in \mathcal{S} \mid \vec{s} \models Init \}$ is the set of initial states, $\mathcal{R} = \{ (\vec{s}, \vec{t}) \mid \vec{s}, \vec{t} \in \mathcal{S}, \exists \vec{l} \in L \cdot Tr(\vec{s}, \vec{l}, \vec{t}') \}$ is a transition relation. Throughout, we write \mathcal{S}' for $\{ \vec{s}' \mid \vec{s} \in \mathcal{S} \}$, and \vec{t}' for $\vec{t}(x')$.

Definition 3. Program $P_1 = \langle V \cup L_1 \cup V', Init_1, Tr_1 \rangle$ is an abstraction of program $P_2 = \langle V \cup L_2 \cup V', Init_2, Tr_2 \rangle$ iff

$$Init_2 \implies Init_1 \quad (\exists \vec{l}_2 \in L_2 \cdot Tr_2(\vec{s}, \vec{l}_2, \vec{s}')) \implies (\exists \vec{l}_1 \in L_1 \cdot Tr_1(\vec{s}, \vec{l}_1, \vec{s}')) \quad (1)$$

An example of abstraction AQ of Q is shown on Fig. 2b-2c. In AQ, non-determinism is introduced on the level of the input parameters.

Definition 4 (cf.[20]). Given two transition systems S and T , a relation $\rho \subseteq \mathcal{S}_S \times \mathcal{S}_T$ is a simulation relation if (1) every state in \mathcal{I}_S is related by ρ to some state in \mathcal{I}_T , and (2) for all states \vec{s}, \vec{s}' and \vec{t} , such that $(\vec{s}, \vec{t}) \in \rho$ and $(\vec{s}, \vec{s}') \in \mathcal{R}_S$ there is some state in \mathcal{S}_T such that $(\vec{t}, \vec{t}') \in \mathcal{R}_T$ and $(\vec{s}', \vec{t}') \in \rho$.

<pre> int P (int a, int b, int c) { int m = b + c; int ret = m + a; return ret; } </pre>	<pre> int Q (int a, int b, int c) { int m = a + b; int n = c + 1; int ret = m - n; return ret; } </pre>	<pre> int AQ (int a, int b, *) { int m = a + b; int n = c + 1; int ret = m - n; return ret; } </pre>
(a) Source program P	(b) Target program Q	(c) Abstraction of Q

Fig. 2: Fragments of the three programs in C

We write, $T_1 \preceq_\rho T_2$, to denote that a transition system T_1 is simulated by a transition system T_2 via a simulation relation ρ . We write $T_1 \preceq T_2$ to indicate existence of a simulation between T_1 and T_2 . We extend the notion of a simulation from transition systems to programs in the usual way: a program P_1 is simulated by a program P_2 iff their corresponding transition systems simulate each other.

Lemma 1. *If P_1 is an abstraction of P_2 then the corresponding transition systems $T_2 \preceq_{id} T_1$, where id is the identity relation.*

Lemma 2. *If $T_1 \preceq T_2$ and $T_2 \preceq T_3$ then $T_1 \preceq T_3$.*

3 From Simulation to Satisfiability

In this section, we show that deciding whether a given relation ρ is a simulation relation is reducible to deciding validity of a $\forall\exists$ -formula. We then show how Skolem functions witnessing the validity of the quantifiers can be used to refine a given simulation relation.

3.1 Deciding Simulation Symbolically

Let $S(\vec{s}, \vec{x}, \vec{s}')^1$ and $T(\vec{t}, \vec{y}, \vec{t}')^1$ be formulas denoting transition relations of two programs, where \vec{s} and \vec{t} , \vec{s}' and \vec{t}' , \vec{x} and \vec{y} are current-state, next-state, and local variables, respectively. Let $\rho(\vec{s}, \vec{t})$ denote a relation between state variables of S and T . For simplicity, we omit the arguments and simply write S , T , and ρ , when the arguments are clear from the context.

Lemma 3. *Given programs S and T , a relation ρ is a simulation relation between S and T iff*

$$\rho(\vec{s}, \vec{t}) \wedge \exists \vec{x} \cdot S(\vec{s}, \vec{x}, \vec{s}') \implies \exists \vec{t}', \vec{y} \cdot T(\vec{t}, \vec{y}, \vec{t}') \wedge \rho(\vec{s}', \vec{t}') \quad (2)$$

The left-hand-side of implication (2) represents the set of all behaviors of S and the set of all matched input conditions. The right-hand-side of (2) represents the existence of a matching behavior in T .

Example 1. Consider two programs P and Q shown on Fig. 2a and Fig. 2b, respectively. Their corresponding transition relations are shown in (3):

$$\begin{aligned} P &\equiv m_P = b_P + c_P \wedge ret_P = m_P + a_P \\ Q &\equiv m_Q = a_Q + b_Q \wedge n_Q = c_Q + 1 \wedge ret_Q = m_Q + n_Q \end{aligned} \quad (3)$$

where the subscript indicates which program the corresponding variables is defined in.

Let \mathcal{C} and \mathcal{N} be relations between current and next-state variables of P and Q , respectively, defined as follows:

$$\mathcal{C} \equiv a_P = a_Q \wedge b_P = b_Q \wedge c_P = c_Q \quad \mathcal{N}' \equiv ret_P = ret_Q \quad (4)$$

Note that in general, unlike in our simplified definition in Section 2, current and next-state variables of a program can differ, requiring us to split the simulation relation into two components.

\mathcal{C} and \mathcal{N} is a simulation relation iff the following formula is valid:

$$\mathcal{C} \wedge (\exists m_P \cdot P) \implies \exists ret_Q, m_Q, n_Q \cdot Q \wedge \mathcal{N} \quad (5)$$

Note that since Q is deterministic, the existential quantifiers in (5) are eliminated trivially by substitution. In our example, (5) simplifies to $0 = 1$, hence \mathcal{C} and \mathcal{N} is *not* a simulation relation between P and Q . \square

3.2 Abstract Simulation

In this section, we show how to check whether a given relation ρ is a simulation relation between a program S and an abstraction αT of a program T . We restrict our attention to existential abstraction [6], although the results extend easily to predicate abstraction as well. Our key result is to show that simulation checking can be done without constructing an abstraction explicitly.

Let T be a transition relation of a program, and U a sub-set of the state-variables of T . An existential abstraction, $\alpha_U^{\exists}(T)$, of T abstracts from T all variables in U . Formally, $\alpha_U^{\exists}(T) \equiv \exists U, U' \cdot T(\vec{t}, \vec{y}, \vec{t}')$, where $U \subseteq \vec{t}$. For example, the program **AQ** in Fig. 2c is an existential abstraction of the program Q in Fig. 2b, where $U = \{c\}$.

Deciding whether a given relation ρ_α is a simulation between a concrete program S and an abstract program αT can be done without computing the abstraction explicitly. Intuitively, the variables that are abstracted away are simply treated as local variables of T .

Lemma 4. *Let $S(\vec{s}, \vec{x}, \vec{s}')$ and $T(\vec{t}, \vec{y}, \vec{t}')$ be two programs. Let $U \subseteq \vec{t}$ be the set of abstracted variables and $\vec{t}_1 = \vec{t} \setminus U$. A relation ρ_α is a simulation relation between S and $\alpha_U^{\exists}(T)$ iff*

$$\rho_\alpha(\vec{s}, \vec{t}_1) \wedge S(\vec{s}, \vec{x}, \vec{s}') \implies \exists \vec{t}_1, \vec{y}, U, U' \cdot T(\vec{t}, \vec{y}, \vec{t}') \wedge \rho_\alpha(\vec{s}', \vec{t}_1) \quad (6)$$

Proof. Immediate from (2) and the definition of existential abstraction. \square

Recall that in Example 1, program P was shown to be not simulated by Q via identity relation. Interestingly, this result is still useful to obtain a valid simulation relation between P and Q by creating implicit abstraction of Q and further refining it. We will demonstrate this 2-steps procedure in Example 2.

Example 2. As a first (*abstraction*) step, we create an implicit abstraction of Q by choosing a state-variable (let it be c) to be existentially quantified. Note that the produced abstraction is equivalent to AQ . Instead of encoding a transition system AQ from scratch (similarly to (3)), we let it be equal to $AQ \equiv \exists c_Q \cdot Q$

Relation \mathcal{C} and \mathcal{N} (disproven to be a simulation relation for P and Q) are abstracted away in correspondence with AQ :

$$\mathcal{C}_\alpha \equiv a_P = a_Q \wedge b_P = b_Q \qquad \mathcal{N}_\alpha \equiv ret_P = ret_Q \qquad (7)$$

\mathcal{C} and \mathcal{N} is a simulation relation between P and AQ iff the following formula is valid:

$$\mathcal{C} \wedge (\exists m_P \cdot P) \implies \exists c_Q, ret_Q, m_Q, n_Q \cdot Q \wedge \mathcal{N} \qquad (8)$$

Clearly, (8) is valid iff there is a Skolem function for the existentially quantified variable c_Q . Note that $sk_{c_Q}(c_P) = -c_P - 1$ is such a function:

$$\mathcal{C} \wedge (\exists m_P \cdot P) \implies (c_Q = -c_P - 1 \implies \exists ret_Q, m_Q, n_Q \cdot Q \wedge \mathcal{N}) \qquad (9)$$

As a second (*refinement*) step, sk_{c_Q} is used to strengthen a simulation relation \mathcal{C} and \mathcal{N} between P and AQ .

$$\mathcal{C}_\alpha^{ext} \equiv a_P = a_Q \wedge b_P = b_Q \wedge c_Q = -c_P - 1 \qquad \mathcal{N}_\alpha^{ext} \equiv ret_P = ret_Q \qquad (10)$$

Note that \mathcal{C}_α^{ext} and \mathcal{N}_α^{ext} is a simulation relation between P and Q. \square

Detailed definition of the Skolem function, its generalization and application to the simulation-relation-checking problem is given in Sect. 3.3.

3.3 Refining Simulation by Skolem Relations

In this section, we show how to use a Skolem relation that is witnessing the validity of the abstract simulation check (6) to refine an abstract simulation relation. We begin with the classical definition of a Skolem function:

Definition 5. *Given a valid formula $\forall x \cdot \exists y \cdot f(x, y)$, a Skolem function for y , $sk_y(x)$ is a function such that $\exists y \cdot f(x, y) \iff f(x, sk_y(x))$.*

We now relax the definition by allowing the relationship between y and x to be an arbitrary relation:

Definition 6. *Given a valid formula $\forall x \cdot \exists y \cdot f(x, y)$, a Skolem relation for y is a relation $Sk_y(x, y)$ such that $\exists y \cdot f(x, y) \iff (Sk_y(x, y) \implies f(x, y))$.*

To see that Def. 6 is a generalization of Def. 5, let $sk_y(x)$ be a Skolem function of y in $\exists y \cdot f(x, y)$. Then, $Sk_y(x, y) \equiv (y = sk_y(x))$ is the corresponding Skolem relation.

$$\left[Sk_y(x, y) \implies f(x, y) \right] \equiv \left[y = sk_y(x) \implies f(x, y) \right] \equiv \left[f(x, sk_y(x)) \right] \quad (11)$$

Clearly, the opposite is not true – not every relation can be represented by a Skolem function.

As shown above, a Skolem relation eliminates an existential quantifier in a valid $\forall\exists$ -formula. In fact, validity of a $\forall\exists$ -formula is equivalent to an existence of an appropriate Skolem function (or relation). We now adapt this to the case of simulation checking.

Theorem 1. *Let $S(\vec{s}, \vec{x}, \vec{s}')$ and $T(\vec{t}, \vec{y}, \vec{t}')$ be two programs such that $S \preceq T$, and $U \subseteq \vec{t}$. Let ρ_α be a simulation relation such that $S \preceq_{\rho_\alpha} \alpha_U^\exists(T)$. Then, there exists a relation $Sk(\vec{s}, U)$ such that (a) $\rho_\alpha \wedge Sk$ is a simulation relation between S and T and (b) Sk is a Skolem relation for U in (6).³*

Recall that by Lemma 3, simulation checking between transition systems S and T is reduced to deciding validity of formula (2). In the next section, we will focus on solving (2) by iterative quantifier elimination and present a generalized algorithm for it.

4 Deciding Validity of $\forall\exists$ -formulas and Extracting Skolem Relations

In this section, we present a novel algorithm, AE-VAL, for deciding validity of $\forall\exists$ -formulas. Without loss of generality, we restrict the input formula to the form $S(\vec{x}) \implies \exists \vec{y} \cdot T(\vec{x}, \vec{y})$, where S has no universal quantifiers, and T is quantifier-free.

4.1 Deciding Validity of $\forall\exists$ -formulas

Our algorithm is based on a notion of Model-Based Projection (MBP), introduced in [15], that under-approximates existential quantification. Let M be a model of a formula $T(\vec{x}, \vec{y})$. Then, $T_{\vec{y}}(\vec{x})$ is an MBP if the following two conditions hold: (a) $M \models T_{\vec{y}}(\vec{x})$ and (b) $T_{\vec{y}}(\vec{x}) \implies \exists \vec{y} \cdot T(\vec{x}, \vec{y})$. That is, the only \vec{x} variables appear in T , $T_{\vec{y}}$, M is a model of $T_{\vec{y}}$, and $T_{\vec{y}}$ is an implicant of T . Furthermore, when \vec{y} and T are fixed, MBP is finite. That is, there are finitely many projections $T_{\vec{y}_1}(\vec{x}), T_{\vec{y}_2}(\vec{x}), \dots, T_{\vec{y}_n}(\vec{x})$ such that $(\exists \vec{y} \cdot T(\vec{x}, \vec{y})) = \bigvee_{i=1}^n T_{\vec{y}_i}(\vec{x})$ for some n . In our implementation, we are using an MBP function from [15] for LRA that is based on Loos-Weispfenning [16] quantifier elimination. Additionally, we assume that for each projection $T_{\vec{y}_i}$ the MBP procedure produces a *local Skolem* relation $\phi_i(\vec{x}, \vec{y})$ such that $\phi_i(\vec{x}, \vec{y}) \implies (T_{\vec{y}_i}(\vec{x}) \implies T(\vec{x}, \vec{y}))$. Local Skolems are a natural by-product of the MBP algorithm in [15]. We write

³ Due to lack of space, the proof is moved to Appendix A.

Algorithm 1: AE-VAL ($S(\vec{x}), \exists \vec{y} \cdot T(\vec{x}, \vec{y})$)

Input: $S(\vec{x}), \exists \vec{y} \cdot T(\vec{x}, \vec{y})$
Output: $res \in \{\text{VALID}, \text{INVALID}\}$ of $S(\vec{x}) \implies \exists \vec{y} \cdot T(\vec{x}, \vec{y})$
Data: Incremental SMT SOLVER, Model M ,
 Model-based projection $T_{\vec{y}}(\vec{x})$, local Skolem relation $\phi_i(\vec{x}, \vec{y})$

```

1 SMTASSERT( $S(\vec{x})$ );
2 while true do
3    $res \leftarrow$  SMTSOLVE();
4   if (ISUNSAT( $res$ )) then return VALID;
5   SMTPOP();
6   SMTADD( $T(\vec{x}, \vec{y})$ );
7    $res \leftarrow$  SMTSOLVE();
8   if (ISUNSAT( $res$ )) then return INVALID;
9    $M \leftarrow$  SMTGETMODEL( $T(\vec{x}, \vec{y})$ );
10   $T_{\vec{y}}, \phi(\vec{x}, \vec{y}) \leftarrow$  GETMBP( $\vec{y}, M, T(\vec{x}, \vec{y})$ );
11  SMTPOP();
12  SMTADD( $\neg T_{\vec{y}}$ );
13 end

```




Fig. 3: Illustration to Example 3 with 3 iterations of AE-VAL ($S, \exists T$ - hexagons, MBPs - ovals, models - points)

(T_i, ϕ_i) \leftarrow GETMBP($\vec{y}, M, T(\vec{x}, \vec{y})$) for an MBP algorithm that takes a formula T , a model M of T and a vector of variables \vec{y} , and returns a projection T_i of T that covers M and the corresponding local Skolem ϕ_i .

AE-VAL is shown in Alg. 1. Given two formulas $S(\vec{x})$ and $\exists \vec{y} \cdot T(\vec{x}, \vec{y})$ it decides the validity of $S(\vec{x}) \implies \exists \vec{y} \cdot T(\vec{x}, \vec{y})$. AE-VAL enumerates the models of $S \wedge T$. In each iteration, it first checks whether there S is non-empty (line 3) and then looks for a model M of $S \wedge T$ (line 9). If M is found, AE-VAL constructs an MBP $T_{\vec{y}}$ of T and M (line 10) and blocks all models contained in $T_{\vec{y}}$ from S (line 12). It iterates until either it discovers that there is a model of S that can not be extended to a model of T (line 8); or all models of $S(\vec{x})$ are blocked (line 4). In the first case, the formula is invalid. In the second, every model of S has been extended to some model of T , and the formula is valid.

Possible three iterations of AE-VAL are depicted graphically in Fig. 3. In the first iteration, AE-VAL selects a model M_1 and generalizes it to an MBP $T_{\vec{y}_1}$. Then, it picks a model M_2 and generalizes it to an MBP $T_{\vec{y}_2}$. Finally, it picks a model M_3 and generalizes it to $T_{\vec{y}_3}$. At this point, all models of S are covered by \vec{y} -free implicants of T , and the algorithm terminates. We demonstrate this further in the following example.

Example 3. Let $\vec{x} \equiv \{a, b\}$, $\vec{y} \equiv \{a', b', c'\}$, and S and T be defined as follows:

$$\begin{aligned}
 S &\equiv (a = b + 2) \\
 T &\equiv (a' = a + b) \wedge (a' = 1 \implies b' = c') \wedge (a' = 2 \implies b' = c' + 1) \wedge \\
 &\quad (a' = 3 \implies b' = c' - 1)
 \end{aligned} \tag{12}$$

We use Φ_i to denote the formula in the SMT context at the beginning of iteration i of AE-VAL. Initially, $\Phi_1 = S$. The first model is $M_1 \equiv \{a = 0, b = -2, a' =$

$-2, b' = 0, c' = 0\}$. $\text{GETMBP}(\vec{y}, M_1, T)$ returns:

$$T_1 \equiv (a + b \neq 2) \wedge (a + b \neq 3) \quad \phi_1 \equiv (a' = a + b) \wedge (b' = c') \quad (13)$$

In the second iteration, $\Phi_2 = \Phi_1 \wedge \neg T_1$, $M_2 \equiv \{a = 2, b = 1, a' = 3, b' = 0, c' = 1\}$, and $\text{GETMBP}(\vec{y}, M_2, T)$ returns:

$$T_2 \equiv (a + b \neq 1) \wedge (a + b \neq 2) \quad \phi_2 \equiv (a' = b + a) \wedge (b' = c' - 1) \quad (14)$$

In the third iteration, $\Phi_3 = \Phi_2 \wedge \neg T_2$, $M_3 \equiv \{a = 2, b = 0, a' = 2, b' = 1, c' = 0\}$, and $\text{GETMBP}(\vec{y}, M_3, T)$ returns:

$$T_3 \equiv (a + b \neq 1) \wedge (a + b \neq 3) \quad \phi_3 \equiv (a' = a + b) \wedge (b' = c' + 1) \quad (15)$$

Since $\Phi_4 = \Phi_3 \wedge \neg T_3$ is UNSAT, EA-VAL returns VALID and terminates. \square

4.2 Extracting Skolem Relation

In the previous section, we have shown an algorithm AE-VAL that decides validity of $S(\vec{x}) \implies \exists \vec{y} \cdot T(\vec{x}, \vec{y})$. As a by-product, it constructs a set of MBPs $\{T_{\vec{y}_i}(\vec{x})\}$ for T and the corresponding local Skolem relations $\{\phi_{\vec{y}_i}(\vec{x}, \vec{y})\}$. In this section, we show how this information can be turned into a Skolem relation $Sk_{\vec{y}}(\vec{x}, \vec{y})$.

Intuitively, $Sk_{\vec{y}}(\vec{x}, \vec{y})$ maps each model of S to a corresponding model of T . However, the local Skolem relation $\phi_{\vec{y}_i}(\vec{x}, \vec{y})$ provides only a partial map (i.e., only for the subset $S(\vec{x}) \wedge T_{\vec{y}_i}(\vec{x})$ of S). Moreover, the local Skolem relations are not disjoint (e.g., see Fig. 3). Thus, to define the Skolem relation Sk , we need to address two issues: (1) we need to find a partitioning $\{I_i\}_{i=1}^n$ of S , and (2) each partition must be associated with an appropriate local Skolem relation.

The constraints on the partitions I_i are as follows. First, a partition I_i must cover all models of $T_{\vec{y}_i}$ that are not already covered by other elements of the partition. Second, it should not include any models that are not contained in $T_{\vec{y}_i}$. Writing these requirements formally, we get the following system of constraints:

$$\begin{cases} S(\vec{x}) \wedge T_{\vec{y}_1}(\vec{x}) \implies I_1(\vec{x}) \\ S(\vec{x}) \wedge T_{\vec{y}_2}(\vec{x}) \wedge \neg T_{\vec{y}_1}(\vec{x}) \implies I_2(\vec{x}) \\ \dots \\ S(\vec{x}) \wedge T_{\vec{y}_n}(\vec{x}) \wedge \neg T_{\vec{y}_1}(\vec{x}) \wedge \neg T_{\vec{y}_2}(\vec{x}) \wedge \dots \wedge \neg T_{\vec{y}_{n-1}}(\vec{x}) \implies I_n(\vec{x}) \\ S(\vec{x}) \wedge I_1(\vec{x}) \wedge \neg T_{\vec{y}_1}(\vec{x}) \implies \perp \\ \dots \\ S(\vec{x}) \wedge I_n(\vec{x}) \wedge \neg T_{\vec{y}_n}(\vec{x}) \implies \perp \end{cases} \quad (16)$$

Note that in (16), S and $\{T_{\vec{y}_i}\}$ are first-order formulas, and $\{I_i\}$ are uninterpreted predicates. The set of constraints corresponds to a set of recursion-free Horn clauses. Thus, we can find an interpretation of the predicates $\{I_i\}$ using a Horn-clause solver. In our implementation, we use the solver of Z3, but other solutions, for example, based on interpolation, are also possible.

We now define the Skolem relation $Sk_{vy}(\vec{x}, \vec{y})$ as follows:

$$Sk_{\vec{y}}(\vec{x}, \vec{y}) \equiv \begin{cases} \phi_{\vec{y}_1}(\vec{x}, \vec{y}) & \text{if } I_1(\vec{x}) \\ \dots & \\ \phi_{\vec{y}_n}(\vec{x}, \vec{y}) & \text{else if } I_n(\vec{x}) \end{cases} \quad (17)$$

The following two theorems show the soundness and completeness of our Skolem relation $Sk_{\vec{y}}(\vec{x}, \vec{y})$. Soundness means that for chosen model \vec{x} , $Sk_{\vec{y}}(\vec{x}, \vec{y})$ satisfies the Def. 6. Completeness means that $Sk_{\vec{y}}(\vec{x}, \vec{y})$ is defined for all models of \vec{x} .

Theorem 2 (Soundness of Skolem Relation). *If the set $\{I_i(\vec{x})\}$ is a solution to (16), and $Sk_{\vec{y}}(\vec{x}, \vec{y})$ is as in (17) then: $S(\vec{x}) \wedge Sk_{\vec{y}}(\vec{x}, \vec{y}) \implies T(\vec{x}, \vec{y})$.*

Proof. Simplifying (17), we need to prove that for every $1 \leq i \leq n$, $S(\vec{x}) \wedge I_i(\vec{x}) \wedge \phi_{\vec{y}_i}(\vec{x}, \vec{y}) \implies T(\vec{x}, \vec{y})$. It is enough to prove that for every i , $S(\vec{x}) \wedge I_i(\vec{x}) \implies T_{\vec{y}_i}(\vec{x})$, which is guaranteed by the last n constraints of (16). \square

Theorem 3 (Completeness of Skolem Relation). *If the set $\{I_i(\vec{x})\}$ is a solution to (16) then $S(\vec{x}) \implies \bigvee_i^n I_i(\vec{x})$.*

Proof. Follows immediately from the first n constraints of (16). \square

Example 4. A partitioning I_1, I_2, I_3 that determines a Skolem relation for Example 3 is: $I_1 \equiv (b \neq 1) \wedge (b \neq 0)$, $I_2 \equiv b \geq 1$, and $I_3 \equiv b = 0$. \square

Constructing a Minimal Skolem relation. Any solution to (16) creates a Skolem relation. But not all Skolem relations are equal. In practice, we often like a Skolem relation that minimizes the number of variables on which each partition depends. For example, in Example 4, we have chosen a partition that only depends on the variable b alone. A simple way to find a minimal solution is to iteratively restrict the number of variables in each partition in (16) until no smaller solution can be found. We leave the problem of finding the minimum partitioning for future work.

5 Simulation-Abstraction-Refinement Loop

In this section, we present our algorithm SIMABS that iteratively constructs and strengthens a simulation relation. While so far we have assumed that a program is represented by a single transition relation, in practice, our algorithms operate on the Cut Point Graph (CPG) [12] representation. A CPG (or a Large Block Encoding (LBE) [4]) is a generalization of a Control Flow Graph (CFG) in which nodes (called *cutpoints*) correspond to loop heads and edges to longest loop free program fragments. In this section, we use CPGs informally and refer the reader to [12] for the formal definition.

The main loop of SIMABS is shown in Alg. 2. The input is two programs S and T . The output is an abstraction αT of T and a simulation relation ρ_α such

Algorithm 2: SIMABS(S, T)

Input: programs S and T ,
abstraction quality metric $\mathcal{Q} : \alpha \rightarrow \{\top, \perp\}$
Output: an abstraction αT ,
and a simulation relation ρ_α , s.t. $S \preceq_{\rho_\alpha} \alpha T$

```
1  $\rho \leftarrow \text{SYNTHESIZE}(S, T, \emptyset)$ ;  
2 while ( $\top$ ) do  
3    $\alpha_{pre} T \leftarrow \alpha T$ ;  
4    $\alpha T, \rho_\alpha \leftarrow \text{FINDABS}(S, T, \rho_\alpha)$ ;  
5   if ( $\alpha T \neq \mathbb{U}$ ) then return  $\mathbb{U}, \emptyset$ ;  
6    $\alpha T, \rho_\alpha \leftarrow \text{EXTEND}(S, \alpha T, \rho_\alpha)$ ;  
7   if ( $\mathcal{Q}(\alpha T) \vee (\alpha_{pre} T = \alpha T)$ ) then return  $\alpha T, \rho_\alpha$ ;
```

Algorithm 3: FINDABS(S, T, ρ)

Input: programs S and T , and a relation ρ
Output: an abstraction αT ,
and a simulation relation ρ_α , s.t. $S \preceq_\alpha \alpha T$

```
1 for each  $(u, v) \in E$  do  
2   if ( $\tau_S(u, v) \not\preceq_\rho \tau_T(u, v)$ ) then  
3      $\alpha T \leftarrow \text{WEAKEN}(T)$ ;  
4     if ( $\alpha T \neq \mathbb{U}$ ) then  
5        $\rho_\alpha \leftarrow \text{SYNTHESIZE}(S, \alpha T, \emptyset)$ ;  
6       return  $\text{FINDABS}(S, \alpha T, \rho_\alpha)$   
7     else return  $\mathbb{U}, \emptyset$ ;  
8 return  $T, \rho$ ;
```

Algorithm 4: EXTEND($S, \alpha T, \rho_\alpha$)

Input: S , abstraction αT , simulation relation ρ_α
Output: abstraction $\alpha^{ext} T$, simulation relation ρ_α^{ext}

```
1  $\rho_\alpha^{ext} \leftarrow \rho_\alpha$ ;  $\alpha^{ext} T \leftarrow \alpha T$ ;  
2  $WL \leftarrow E$ ;  
3 while ( $WL \neq \emptyset$ ) do  
4    $(u, v) \leftarrow \text{GETWTO SMALLESTEDGE}(WL)$ ;  
5    $WL \leftarrow WL \setminus \{(u, v)\}$ ;  
6   if ( $\tau_S(u, v) \preceq_{\rho_\alpha^{ext}} \tau_{\alpha^{ext} T}(u, v)$ ) then  
7      $Sk \leftarrow \text{EXTRACTSKOLEM}(u, v, \rho_\alpha^{ext})$ ;  
8      $\alpha^{ext} T \leftarrow \text{STRENGTHEN}(\alpha T, Sk)$ ;  
9      $\rho_\alpha^{ext} \leftarrow \text{SYNTHESIZE}(S, \alpha^{ext} T, Sk)$ ;  
10     $WL \leftarrow WL \cup \{(v, x) \in E \mid x \in \text{CP}\} \cup \{(y, v) \in E \mid y \in \text{CP}\}$ ;  
11  else return  $\alpha T, \rho_\alpha$ ;  
12 return  $\alpha^{ext} T, \rho_\alpha^{ext}$ 
```

that $S \preceq_{\rho_\alpha} T$, if such an abstraction exists. In the presentation, we assume that S and T share the set of cutpoints CP, but might have different interpretation of control-flow edges by loop-free program fragments. However, our algorithm extends to the general case where S and T have different cutpoints as well.

SIMABS begins by guessing an initial relation ρ using SYNTHESIZE. The initial guess can be an arbitrary relation between the CPGs of S and T . In our implementation, for every cutpoint, we take ρ to be the identity relation between the live variables of S and T at that cutpoint, that have identical names. Then, SIMABS uses FINDABS to check whether there exists an abstraction of ρ (including ρ itself) that is a simulation relation between S and a corresponding abstraction αT of T . If this succeeds, the method EXTEND is used to refine the abstraction αT and synthesize a new simulation relation ρ_α . This process continues until the abstraction is satisfied by some quality metric \mathcal{Q} (line 7) (for example, if it is sufficient to (dis-)prove some safety property of T) or it is equivalent to an abstraction produced in the previous iteration (line 3). Otherwise, SIMABS goes into the next iteration and finds another abstraction. If an abstraction cannot be constructed (we use a shortcut \mathbb{U} for this), SIMABS terminates with a negative result (line 5).

FINDABS (shown in Alg. 3) is given S , T and ρ . It traverses the set of CPG edges E , common to S and T (line 1) in the Weak Topological Ordering (WTO) [5] (in which inner loops are traversed before outer loops). For each edge $(u, v) \in E$, FINDABS checks (on line 2) whether ρ is a simulation relation between the loop-free program fragment $\tau_S(u, v)$ labeling the (u, v) -edge in S

and the corresponding loop-free program fragment $\tau_T(u, v)$ labeling the (u, v) -edge in T . If the check succeeds for all edges, the ρ is a simulation relation between S and T , and it is returned to the user. Otherwise, FINDABS chooses an abstraction αT of T using the method WEAKEN, synthesizes a new ρ_α and repeats the check for S , αT , and ρ_α . WEAKEN introduces non-determinism to the interpretation of control-flow edges of T . In our implementation, the most successful implementation of WEAKEN is the existential abstraction of all state-variables that are live at the source and destination of the edge (u, v) for which the simulation check has failed. For each iteration of SIMABS, FINDABS is given a concrete program T , and always constructs a new abstraction from scratch.

EXTEND (shown in Alg. 4) is given S , αT , and ρ_α and constructs a refinement ρ_α^{ext} of simulation relation ρ_α and corresponding refinement $\alpha^{ext}T$ of αT . EXTEND maintains a work-list WL of control-flow edges to be processed. Initially, WL is populated with all the CPG edges E (line 2). In each iteration, EXTEND strengthens αT and synthesizes a new ρ_α by strengthening the old ρ_α by a Skolem relation Sk . Sk is produced from the proof of the simulation between fragments $\tau_S(u, v)$ and $\tau_{\alpha T}(u, v)$ of the smallest element (u, v) of E according to the WTO (line 6). Finally, EXTEND updates WL (line 10) with the edges that share the next-state variables that have appeared in Sk and iterates until WL is empty. If in some iteration, a strengthening is impossible, EXTEND returns the last successful values for ρ_α^{ext} and $\alpha^{ext}T$.

Recall programs P and Q from Fig. 2a and Fig. 2b. In Examples 1 and 2, we proved that $P \preceq Q$. In the following , we show how SYMABS automates this.

Example 5. First, SYMABS is given P and Q ; SYNTHESIZE guesses \mathcal{C} and \mathcal{N} as in (4). Then FINDABS disproves that \mathcal{C} and \mathcal{N} is a simulation relation, but chooses an implicit abstraction of Q (equivalent to AQ on Fig. 2c), and constructs \mathcal{C}_α and \mathcal{N}_α , as in (7), for which the simulation-relation-formula (8) for P and AQ is valid. Finally, EXTEND extracts Skolem relation and uses it to create \mathcal{C}_α^{ext} and \mathcal{N}_α^{ext} , as in (10) and confirms that it is a simulation relation for P and Q . \square

6 Evaluation

We have implemented SIMABS in the UFO framework, and evaluated it on the Software Verification Competition (SVCOMP'14) benchmarks and `instcombine` and `simplifycfg` optimizations of LLVM. The `instcombine` performs local arithmetic optimizations (e.g., replacing `a = 1 - 1` by `a = 0`). The `simplifycfg` performs dead code elimination and basic block merging (e.g., replacing `if (true) {a++;} else {a--;} by a++`). The combination of these optimizations provides more aggressive optimizations.

For each source benchmark (S) (300 - 5000 lines of source code), we created a target optimization (T) by two applications of (`instcombine + simplifycfg`). We applied SIMABS to discover two simulations: $S \preceq T$ and $T \preceq S$. Out of all benchmarks, we chose 157 for which SIMABS terminates within 5 minutes. We

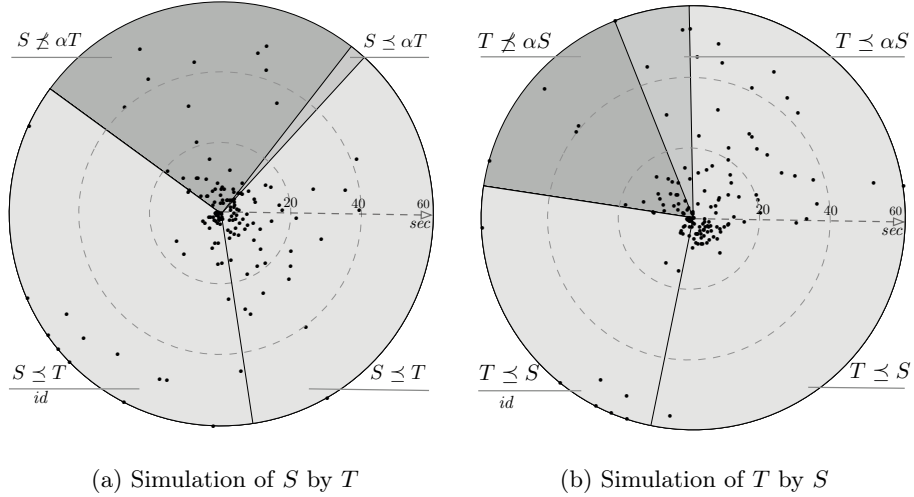


Fig. 4: Pie chart and running times in spherical coordinate system.

present the results in two diagrams in Fig. 4. Full results are available at <http://www.inf.usi.ch/phd/fedyukovich/simabs.pdf>⁴.

Each diagram is a pie chart and a collection of SIMABS execution times for each benchmark in the spherical coordinate system. The pie chart in Fig. 4a represents a proportion of four main classes of SIMABS results: whether (a-b) T simulates S (in (a), via identity relation), (c) T does not simulate S , but some abstraction αT does, (d) T does not simulate S and we did not find an abstraction αT that does. Each dot represents a runtime of SIMABS on a single benchmark. It is placed in one of the segments (a)–(d) with respect to the outcome, and is assigned the unique polar angle and the radial distance to represent time in seconds. For example, a benchmark on which $S \preceq_{id} T$ solved in 20s is placed in quadrant (a) in a distance 20 from the center. Closer to the center means faster. Runs that took longer than 60s are placed on the boundary. Fig. 4b is structured similarly, but with inverse order of S and T .

The experiment confirms our intuition that the original program more often simulates the optimal one than vice versa. According to Fig. 4, in 115 cases $S \preceq \alpha T$, and in 122 cases $T \preceq \alpha S$; in 40 cases $S \not\preceq \alpha T$, and in 27 cases $T \not\preceq \alpha S$.

The experiment confirms that the use of our novel simulation-abstraction-refinement loop by SIMABS is necessary in majority of cases. There is a relatively small subset of benchmarks (59 of $S \preceq T$ and 39 of $T \preceq S$), in which the identity relation, guessed at the first iteration of SIMABS is already a simulation relation (but still, in many cases it took more than a minute to establish that), while in the rest of the cases, SIMABS goes into the simulation-abstraction-refinement loop and successfully terminates (faster than a minute).

⁴ For convenience, we have included the table of results in an appendix.

7 Related Work

The algebraic notion of simulation relation between programs dates back to Milner [18]. The approach indirectly refers to simulation relation between S and the abstraction of T , by introducing *weak simulation*, and to simulation relation between S and T as *strong simulation*. While being purely theoretic, this work does not consider a practical application of using weak simulations in order to construct strong simulation.

Translation Validation for optimizing compiler [21] checks simulation relation between a program and its GCC-optimization. As a secondary result, that work proposes a simple heuristic to construct simulation relation, restricted to specific optimizations. Checking is done on the level of control-flow graphs and takes into account all program variables. In contrast, our SIMABS algorithm is able to find simulation relation independently on an optimizer.

Classical approach to check that a relation is a simulation relation is by *game-theoretic* approach, in which the state space of the source and the target is traversed by the evader and a pursuer solvers. For instance, [13] applies it to prove simulation relation between infinite graphs. In our setting, this result can be used to extend SIMABS to deal with programs with different CPGs.

The need of eliminating quantifiers by a method AE-VAL makes our approach similar to template-based synthesis [24, 3, 2, 17]. The goal of the approach is to synthesize an arbitrary program that fulfills a given specification represented by a template. While instantiating existential quantifiers, synthesis is filling placeholders in the predefined template formula. While discovering a Skolem relation on the top of valid simulation-relation-checking formula, we also perform a synthesis, but do not require any template for it.

Apart of discovering simulation between programs, there exist another ways to prove their equivalence. For example, rather practical solution to check equivalence between a Verilog circuit and C program was established in [7]. It is based on translation of both programs into quantifier-free propositional formula, satisfiable iff the circuit and the program disagree.

References

1. A. Albarghouthi, A. Gurfinkel, and M. Chechik. UFO: A Framework for Abstraction- and Interpolation-Based Software Verification. In *CAV*, 2012.
2. R. Alur, R. Bodík, G. Juniwal, M. M. K. Martin, M. Raghothaman, S. A. Seshia, R. Singh, A. Solar-Lezama, E. Torlak, and A. Udupa. Syntax-guided synthesis. In *FMCAD*, pages 1–17, 2013.
3. T. A. Beyene, C. Popeea, and A. Rybalchenko. Solving existentially quantified horn clauses. In *CAV*, pages 869–882, 2013.
4. D. Beyer, A. Cimatti, A. Griggio, M. E. Keremoglu, and R. Sebastiani. Software Model Checking via Large-Block Encoding. In *FMCAD*, pages 25–32, 2009.
5. F. A. Bourdoncle. Efficient Chaotic Iteration Strategies with Widenings. In *Proc. of FMPA '93*, LNCS, pages 128–141, 1993.
6. E. M. Clarke, O. Grumberg, and D. E. Long. Model checking and abstraction. *ACM Trans. Program. Lang. Syst.*, 16(5):1512–1542, 1994.

7. E. M. Clarke, D. Kroening, and K. Yorav. Behavioral consistency of c and verilog programs using bounded model checking. In *DAC*, pages 368–371, 2003.
8. L. M. de Moura and N. Bjørner. Z3: An Efficient SMT Solver. In *TACAS*, pages 337–340, 2008.
9. G. Fedyukovich, A. Gurfinkel, and N. Sharygina. Incremental verification of compiler optimizations. In *NFM*, volume 8430 of *LNCS*. Springer, 2014.
10. G. Fedyukovich, O. Sery, and N. Sharygina. eVolCheck: Incremental Upgrade Checker for C. In *Proc. of TACAS'13*, volume 7795 of *LNCS*, pages 292–307, 2013.
11. B. Godlin and O. Strichman. Regression verification. In *DAC*, 2009.
12. A. Gurfinkel, S. Chaki, and S. Sapra. Efficient Predicate Abstraction of Program Summaries. In *NASA Formal Methods*, pages 131–145, 2011.
13. M. R. Henzinger, T. A. Henzinger, and P. W. Kopke. Computing simulations on finite and infinite graphs. In *FOCS*, pages 453–462, 1995.
14. M. Kawaguchi, S. K. Lahiri, and H. Rebelo. Conditional equivalence. Technical Report MSR-TR-2010-119, Microsoft Research, 2010.
15. A. Komuravelli, A. Gurfinkel, and S. Chaki. SMT-Based Model Checking for Recursive Programs. In *CAV*, pages 17–34, 2014.
16. R. Loos and V. Weispfenning. Applying linear quantifier elimination. *Comput. J.*, 36(5):450–462, 1993.
17. R. Madhavan and V. Kuncak. Symbolic resource bound inference for functional programs. In *Computer Aided Verification (CAV)*, 2014.
18. R. Milner. An algebraic definition of simulation between programs. In *IJCAI*, pages 481–489, 1971.
19. K. S. Namjoshi. Lifting Temporal Proofs through Abstractions. In *VMCAI*, volume 2575 of *Lecture Notes in Computer Science*, pages 174–188. Springer, 2003.
20. K. S. Namjoshi and L. D. Zuck. Witnessing program transformations. In *SAS*, 2013.
21. G. C. Necula. Translation validation for an optimizing compiler. In *PLDI*, 2000.
22. S. Person, M. B. Dwyer, S. G. Elbaum, and C. S. Pasareanu. Differential symbolic execution. In *Foundations of SW Engineering (FSE '08)*, pages 226–237, 2008.
23. A.-D. Phan, N. Bjørner, and D. Monniaux. Anatomy of alternating quantifier satisfiability (work in progress). In *SMT*, pages 120–130, 2012.
24. S. Srivastava, S. Gulwani, and J. S. Foster. From program verification to program synthesis. In *POPL*, pages 313–326, 2010.

A The proof of Theorem 1.

Theorem 1. *Let $S(\vec{s}, \vec{x}, \vec{s}')$ and $T(\vec{t}, \vec{y}, \vec{t}')$ be two programs such that $S \preceq T$, and $U \subseteq \vec{t}$. Let ρ_α be a simulation relation such that $S \preceq_{\rho_\alpha} \alpha_{\vec{U}}^{\exists}(T)$. Then, there exists a relation $Sk(\vec{s}, U)$ such that (a) $\rho_\alpha \wedge Sk$ is a simulation relation between S and T and (b) Sk is a Skolem relation for U in (6).*

Proof. Let $\vec{t} \equiv \vec{t}_1 \cup U$ and $\vec{t}' \equiv \vec{t}'_1 \cup U'$. Since $S \preceq T$, there exist relation ρ such that $S \preceq_\rho T$ and $\rho \implies \rho_\alpha$. Let Sk be a relation over \vec{s}, \vec{s}', U and U' , such that

$$\rho(\vec{s}, \vec{t}) \equiv \rho_\alpha(\vec{s}, \vec{t}_1) \wedge Sk(\vec{s}, U) \quad \rho(\vec{s}', \vec{t}') \equiv \rho_\alpha(\vec{s}', \vec{t}'_1) \wedge Sk(\vec{s}', U') \quad (18)$$

From (2) and (18), we get

$$\begin{aligned} \rho_\alpha(\vec{s}, \vec{t}_1) \wedge Sk(\vec{s}, U) \wedge S(\vec{s}, \vec{s}') \implies \\ \exists \vec{t}'_1, U', \vec{y} \cdot T(\vec{t}_1, U, \vec{y}, \vec{t}'_1, U') \wedge \rho_\alpha(\vec{s}', \vec{t}'_1) \wedge Sk(\vec{s}', U') \end{aligned} \quad (19)$$

In (19), Sk is witnessing a Skolem relation for U . By Def. 6, we get:

$$\begin{aligned} \exists U \cdot \rho_\alpha(\vec{s}, \vec{t}_1) \wedge S(\vec{s}, \vec{s}') \implies \\ \exists \vec{t}'_1, U', \vec{y} \cdot T(\vec{t}_1, U, \vec{y}, \vec{t}'_1, U') \wedge \rho_\alpha(\vec{s}', \vec{t}'_1) \wedge Sk(\vec{s}', U') \end{aligned} \quad (20)$$

Since (20) is valid then a weaker formula (21) is also valid.

$$\exists U \cdot \rho_\alpha(\vec{s}, \vec{t}_1) \wedge S(\vec{s}, \vec{s}') \implies \exists \vec{t}'_1, U, U', \vec{y} \cdot T(\vec{t}_1, U, \vec{y}, \vec{t}'_1, U') \wedge \rho_\alpha(\vec{s}', \vec{t}'_1) \quad (21)$$

Notably, (21) is a simulation-checking-formula (6) for S and $\alpha_{\vec{U}}^{\exists}(T)$. It means, the chosen Sk is also a Skolem relation for (6). \square

B Concrete data

Tables 1, 2, 3 gather statistics for all 3 cases ($S \preceq T$, $S \preceq \alpha T$, and $S \not\preceq \alpha T$) respectively.

name	#Vars	#Nondet		Time	cutpoints
		beg	end		
locks/test_locks_14_unsafe.o3	5	0	0	4.34	0/51
locks/test_locks_15_unsafe.o3	5	0	0	4.32	0/51
locks/test_locks_15_unsafe.o0	4	0	0	0.18	0/5
locks/test_locks_14_unsafe.o0	4	0	0	0.18	0/5
mtdrivers/g10kperf_unsafe.1.cil.o3	5	0	0	1.28	0/5
mtdrivers/g10kperf_unsafe.1.cil.o0	36	0	0	152.23	0/11
mtdrivers/hdfltr_unsafe.1.cil.o3	7	0	0	31.8	0/7
product-lines/misepump_speac3_product0_unsafe.cil.o3	4	0	0	0.84	0/3
product-lines/misepump_speac3_product18_unsafe.cil.o3	4	0	0	0.61	0/3
product-lines/misepump_speac3_product29_unsafe.cil.o0	4	0	0	0.98	0/3
product-lines/misepump_speac3_product05_unsafe.cil.o0	4	0	0	0.58	0/3
product-lines/misepump_speac3_product13_unsafe.cil.o0	4	0	0	0.61	0/3
product-lines/misepump_speac3_product24_unsafe.cil.o0	4	0	0	0.88	0/3
product-lines/misepump_speac3_product08_unsafe.cil.o3	4	0	0	0.56	0/3
product-lines/misepump_speac3_product50_unsafe.cil.o0	4	0	0	0.57	0/3
product-lines/misepump_speac3_product06_unsafe.cil.o0	4	0	0	0.6	0/3
product-lines/misepump_speac3_product26_unsafe.cil.o3	4	0	0	0.6	0/3
product-lines/misepump_speac3_product10_unsafe.cil.o3	4	0	0	0.59	0/3
product-lines/misepump_speac3_product17_unsafe.cil.o0	4	0	0	0.6	0/3
product-lines/misepump_speac3_product32_unsafe.cil.o3	4	0	0	0.92	0/3
product-lines/misepump_speac3_product03_unsafe.cil.o0	4	0	0	0.87	0/3
product-lines/misepump_speac3_product18_unsafe.cil.o0	4	0	0	0.59	0/3
product-lines/misepump_speac3_product11_unsafe.cil.o3	4	0	0	0.85	0/3
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-usb-storage-usb-storage.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	17	0	0	4.22	0/8
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-media-dvb-dvb-usb-dvb-usb-vp7045.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	64	0	0	4.06	0/56
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-scsi-libfc-libfc.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	20	0	0	20.73	0/21
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-mfd-tpm600.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	17	0	0	3.38	0/12
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-staging-keucr-keucr.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	12	0	0	5.48	0/5
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-iscsi-kernelcap.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	15	0	0	5.91	0/5
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-media-dvb-dvb-usb-dvb-usb-vp7045.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	231	0	0	111.35	0/65
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-mfd-tpm600.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	6	0	0	1.34	0/8
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-iscsi-kernelcap.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	16	0	0	5.46	0/5
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-media-dvb-dvb-usb-dvb-usb-vp7045.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	70	0	0	19.85	0/32
ldv-linux-3.4/32_f_cilled_unsafe_const.o3.linux-32-1-drivers-media-dvb-dvb-usb-dvb-usb-vp7045.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	162	0	0	14.82	0/117
system/transmitter_02_unsafe.cil.o3	28	0	0	12.49	0/7
system/token_ring_02_saf.cil.o3	37	0	0	23.80	0/7
system/pc_fifo_1_saf.cil.o3	56	0	0	26.79	0/11
system/pc_fifo_2_saf.cil.o3	105	0	0	34.46	0/13
system/transmitter_01_unsafe.cil.o3	22	0	0	1.93	0/12
system/token_ring_01_unsafe.cil.o0	56	0	0	80.07	0/7
system/transmitter_01_unsafe.cil.o0	39	0	0	22.3	0/7
system/bist_cel_saf.cil.o0	35	0	0	105.95	0/7
system/pc_fifo_1_unsafe.cil.o3	56	0	0	16.13	0/11
system/token_ring_02_unsafe.cil.o3	37	0	0	25.85	0/7
system/pc_fifo_2_unsafe.cil.o3	105	0	0	32.4	0/14
system/transmitter_03_unsafe.cil.o3	38	0	0	119.92	0/7
system/numul_unsafe.cil.o3	57	0	0	7.62	0/11
system/token_ring_01_saf.cil.o3	35	0	0	3.74	0/12
system/token_ring_01_unsafe.cil.o3	27	0	0	1.8	0/7
system/token_ring_01_saf.cil.o0	56	0	0	88.99	0/7

Table 1: Full cycle of SIMABS, $S \preceq T$

name	#Vars	#Nondet	Vars breg abs	Time	cutpoints	
ndrivers/flipgy_safe_2.cil.o3	7	1	1	0.126	0/5	
ndrivers/flipgy_umsafe_1.cil.o0	10	1	1	0.5722	0/7	
ndrivers/flipgy_umsafe_1.cil.o3	3	1	1	0.14	0/5	
product-lines/misepump_specs2_product47_umsafe.cil.o3	8	0	4	0.1433	1/3	
product-lines/misepump_specs1_product49_umsafe.cil.o0	6	0	3	0.581	1/3	
product-lines/misepump_specs3_product46_umsafe.cil.o3	8	0	4	0.2981	1/3	
product-lines/misepump_specs3_product48_umsafe.cil.o3	8	0	4	0.3229	1/3	
product-lines/misepump_specs4_product47_umsafe.cil.o3	6	0	3	0.404	1/3	
product-lines/misepump_specs5_product43_umsafe.cil.o0	6	0	3	0.609	1/3	
product-lines/misepump_specs2_product48_umsafe.cil.o0	8	0	4	0.1009	1/3	
product-lines/misepump_specs1_product45_umsafe.cil.o0	6	0	3	0.403	1/3	
product-lines/misepump_specs4_product43_umsafe.cil.o3	6	0	3	0.412	1/3	
product-lines/misepump_specs1_product39_umsafe.cil.o0	8	0	4	0.1813	1/3	
product-lines/misepump_specs4_product38_umsafe.cil.o0	6	0	3	0.606	1/3	
product-lines/misepump_specs4_product40_umsafe.cil.o0	8	0	4	0.5617	1/3	
product-lines/misepump_specs1_product38_umsafe.cil.o0	6	0	3	0.689	1/3	
product-lines/misepump_specs5_product46_umsafe.cil.o3	6	0	3	0.512	1/3	
product-lines/misepump_specs1_product40_umsafe.cil.o3	6	0	3	0.397	1/3	
product-lines/misepump_specs4_product35_umsafe.cil.o3	6	0	3	0.309	1/3	
product-lines/misepump_specs2_product43_umsafe.cil.o0	10	0	5	0.258	1/3	
product-lines/misepump_specs2_product49_umsafe.cil.o0	10	0	5	0.332	1/3	
product-lines/misepump_specs5_product44_umsafe.cil.o0	8	0	4	0.5207	1/3	
product-lines/misepump_specs4_product45_umsafe.cil.o0	8	0	4	0.1944	1/3	
product-lines/misepump_specs2_product38_umsafe.cil.o0	8	0	4	0.535	1/3	
product-lines/misepump_specs1_product44_umsafe.cil.o0	8	0	4	0.1609	1/3	
product-lines/misepump_specs1_product46_umsafe.cil.o0	8	0	4	0.3402	1/3	
product-lines/misepump_specs1_product44_umsafe.cil.o0	8	0	4	0.68	1/3	
product-lines/misepump_specs2_product41_umsafe.cil.o0	10	0	5	0.3864	1/3	
product-lines/misepump_specs1_product48_umsafe.cil.o3	8	0	4	0.1785	1/3	
product-lines/misepump_specs5_product42_umsafe.cil.o3	6	0	3	0.539	1/3	
product-lines/misepump_specs5_product42_umsafe.cil.o0	8	0	4	0.2823	1/3	
product-lines/misepump_specs1_product36_umsafe.cil.o0	8	0	4	0.1211	1/3	
product-lines/misepump_specs2_product39_umsafe.cil.o3	10	0	5	0.109	1/3	
product-lines/misepump_specs2_product40_umsafe.cil.o0	6	0	3	0.617	1/3	
product-lines/misepump_specs5_product38_umsafe.cil.o3	6	0	3	0.272	1/3	
product-lines/misepump_specs3_product45_umsafe.cil.o0	6	0	3	0.618	1/3	
product-lines/misepump_specs4_product41_umsafe.cil.o0	6	0	3	0.1077	1/3	
product-lines/misepump_specs4_product44_umsafe.cil.o3	6	0	3	0.382	1/3	
product-lines/misepump_specs2_product42_umsafe.cil.o3	8	0	4	0.627	1/3	
product-lines/misepump_specs1_product43_umsafe.cil.o3	8	0	4	0.314	1/3	
product-lines/misepump_specs2_product48_umsafe.cil.o3	10	0	5	0.2422	1/3	
product-lines/misepump_specs2_product44_umsafe.cil.o0	8	0	4	0.54	1/3	
product-lines/misepump_specs4_product41_umsafe.cil.o3	6	0	3	0.508	1/3	
product-lines/misepump_specs1_product37_umsafe.cil.o0	6	0	3	0.672	1/3	
product-lines/misepump_specs5_product48_umsafe.cil.o0	6	0	3	0.819	1/3	
product-lines/misepump_specs5_product35_umsafe.cil.o0	8	0	4	0.1207	1/3	
product-lines/misepump_specs2_product48_umsafe.cil.o3	8	0	4	0.672	1/3	
product-lines/misepump_specs1_product36_umsafe.cil.o0	6	0	3	0.325	1/3	
product-lines/misepump_specs3_product47_umsafe.cil.o0	4	2	2	0.116	0/3	
product-lines/misepump_specs2_product43_umsafe.cil.o0	8	0	4	0.853	1/3	
product-lines/misepump_specs1_product45_umsafe.cil.o0	8	0	4	0.2743	1/3	
product-lines/misepump_specs2_product46_umsafe.cil.o3	6	0	3	0.514	1/3	
product-lines/misepump_specs1_product38_umsafe.cil.o3	6	0	3	0.277	1/3	
product-lines/misepump_specs1_product48_umsafe.cil.o3	6	0	3	0.422	1/3	
product-lines/misepump_specs3_product39_umsafe.cil.o3	8	0	4	0.1471	1/3	
product-lines/misepump_specs4_product43_umsafe.cil.o3	6	0	3	0.393	1/3	
product-lines/misepump_specs4_product46_umsafe.cil.o0	6	0	3	0.739	1/3	
product-lines/misepump_specs1_product40_umsafe.cil.o0	8	0	4	0.3522	1/3	
product-lines/misepump_specs1_product42_umsafe.cil.o0	6	0	3	0.384	1/3	
product-lines/misepump_specs1_product43_umsafe.cil.o0	8	0	4	0.1779	1/3	
product-lines/misepump_specs3_product34_umsafe.cil.o3	6	0	3	0.505	1/3	
product-lines/misepump_specs4_product43_umsafe.cil.o3	8	0	4	0.1624	1/3	
product-lines/misepump_specs1_product40_umsafe.cil.o0	6	0	3	0.642	1/3	
product-lines/misepump_specs3_product40_umsafe.cil.o0	8	0	4	0.4676	1/3	
product-lines/misepump_specs4_product42_umsafe.cil.o0	8	0	4	0.3049	1/3	
product-lines/misepump_specs5_product46_umsafe.cil.o3	8	0	4	0.3091	1/3	
product-lines/misepump_specs5_product38_umsafe.cil.o3	6	0	3	0.288	1/3	
product-lines/misepump_specs1_product44_umsafe.cil.o3	8	0	4	0.1713	1/3	
product-lines/misepump_specs3_product47_umsafe.cil.o0	8	0	4	0.4321	1/3	
product-lines/misepump_specs4_product40_umsafe.cil.o3	8	0	4	0.2321	1/3	
product-lines/misepump_specs3_product42_umsafe.cil.o3	6	0	3	0.733	1/3	
product-lines/misepump_specs3_product44_umsafe.cil.o3	8	0	4	0.1183	1/3	
product-lines/misepump_specs5_product38_umsafe.cil.o0	6	0	3	0.419	1/3	
product-lines/misepump_specs3_product45_umsafe.cil.o0	6	0	3	0.805	1/3	
product-lines/misepump_specs2_product43_umsafe.cil.o0	10	0	5	0.4737	1/3	
product-lines/misepump_specs3_product46_umsafe.cil.o0	6	0	3	0.888	1/3	
product-lines/misepump_specs1_product47_umsafe.cil.o3	6	0	3	0.288	1/3	
product-lines/misepump_specs2_product47_umsafe.cil.o3	8	0	4	0.701	1/3	
product-lines/misepump_specs1_product36_umsafe.cil.o3	8	0	4	0.143	1/3	
product-lines/misepump_specs2_product46_umsafe.cil.o3	8	0	4	0.414	1/3	
product-lines/misepump_specs1_product38_umsafe.cil.o0	8	0	4	0.1173	1/3	
ldv-linux-3.4/32_7_cilled_umsafe_const_ok_linux-32-1-drivers-usb-storage-usb-storage.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	30	1	1	0.823	0/8	
ldv-linux-3.4/32_7_cilled_umsafe_const_ok_linux-32-1-drivers-usb-storage-usb-storage.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	25	1	21	14	1.4	2/9
ldv-linux-3.4/32_7_cilled_umsafe_const_ok_linux-32-1-drivers-usb-storage-usb-storage.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	18	0	15	6	1.92	2/9
ldv-linux-3.4/32_1_cilled_umsafe_ok_ponder_linux-3.4-32-1-drivers-media-dvb-dvb-usb-dvb-usb-as027.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	19	0	15	12	15.23	1/5
ldv-linux-3.4/32_1_cilled_umsafe_ok_ponder_linux-3.4-32-1-drivers-media-dvb-dvb-usb-dvb-usb-pct482.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	85	0	60	62	49.18	1/15
ldv-linux-3.4/32_1_cilled_umsafe_ok_ponder_linux-3.4-32-1-drivers-media-dvb-dvb-usb-dvb-usb-pct482.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	39	0	21	19	63.73	1/15
ldv-linux-3.4/32_7_cilled_umsafe_const_ok_linux-32-1-drivers-etagging-kesir-kesir.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	25	1	1	0.603	0/5	
system/c_hifo_2_umsafe.cil.o0	70	6	6	3	59.9	0/14
system/c_hifo_2_umsafe.cil.o3	70	6	6	3	59.14	0/15
system/kundul_umsafe.cil.o0	40	10	10	4	13.14	0/7

Table 2: Full cycle of SIMABS, $S \preceq \alpha T$

name	#Vars	#Nondet	Vars	Time	cutpoints	
		beg	end			
ndrivers/diskperf_safe.i.cil.o3	5	0	0	1.41	1/5	
product-lines/minepump_spec3_product60_unsafe.cil.o0	6	0	3	9.52	2/3	
product-lines/minepump_spec2_product61_safe.cil.o3	8	0	4	5.42	2/3	
product-lines/minepump_spec1_product64_unsafe.cil.o3	6	0	3	4.38	2/3	
product-lines/minepump_spec3_product62_unsafe.cil.o0	8	0	4	55.3	2/3	
product-lines/minepump_spec3_product61_unsafe.cil.o0	8	0	4	39.18	2/3	
product-lines/minepump_spec4_product60_safe.cil.o0	6	0	3	3.27	1/3	
product-lines/minepump_spec3_product64_safe.cil.o0	6	0	3	11.96	2/3	
product-lines/minepump_spec3_product67_safe.cil.o0	6	0	3	10.95	2/3	
product-lines/minepump_spec3_product69_safe.cil.o3	6	0	3	6.41	2/3	
product-lines/minepump_spec3_product60_safe.cil.o3	6	0	3	6.48	2/3	
product-lines/minepump_spec3_product60_unsafe.cil.o3	8	0	4	39.74	2/3	
product-lines/minepump_spec4_product67_safe.cil.o3	6	0	3	3.12	1/3	
product-lines/minepump_spec1_product63_safe.cil.o3	8	0	4	12.43	2/3	
product-lines/minepump_spec3_product62_safe.cil.o0	6	0	3	13.34	2/3	
product-lines/minepump_spec2_product64_safe.cil.o3	8	0	4	5.42	2/3	
product-lines/minepump_spec1_product65_unsafe.cil.o3	8	0	4	11.84	2/3	
product-lines/minepump_spec3_product68_unsafe.cil.o3	6	0	3	7.88	2/3	
lhw-linux-3.4/32_f_cilled_unsafe_const_of_linux-32_1-drivers-usb-chips-cfi_omdset_0001.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	236	0	146	145	21.11	1/121
lhw-linux-3.4/32_f_cilled_unsafe_const_of_linux-32_1-drivers-input-mouse-synaptics_usb.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	22	1	18	10	8.51	9/10
lhw-linux-3.4/32_1_cilled_safe_of_omdset_linux-3.4-32_1-drivers-media-dvb-drv-usb-az6007.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	43	0	10	7	2.97	2/27
lhw-linux-3.4/32_f_cilled_unsafe_const_of_linux-32_1-drivers-input-mouse-synaptics_usb.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	18	0	13	4	9.34	3/10
lhw-linux-3.4/32_1_cilled_safe_of_omdset_linux-3.4-32_1-drivers-media-dvb-drv-usb-az6007.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o0	87	0	10	10	15.5	1/84
lhw-linux-3.4/32_1_cilled_safe_of_omdset_linux-3.4-32_1-drivers-media-dvb-drv-usb-az6007.ko-ldv_main0_sequence_infinite_withcheck_stateful.cil.out.o3	47	0	10	10	14.15	1/13
system/pc_efifo_1_safe.cil.o0	43	3	20	1	4.9	6/12
system/pc_efifo_1_unsafe.cil.o0	43	3	20	2	3.4	7/13

Table 3: Full cycle of SIMABS, $S \not\leq \alpha T$