

Predicate Generation for Learning-Based Quantifier-Free Loop Invariant Inference*

Yungbum Jung¹, Wonchan Lee¹, Bow-Yaw Wang², and Kwangkuen Yi¹

¹ Seoul National University

² INRIA and Academia Sinica

Abstract. We address the predicate generation problem in the context of loop invariant inference. Motivated by the interpolation-based abstraction refinement technique, we apply the interpolation theorem to synthesize predicates implicitly implied by program texts. Our technique is able to improve the effectiveness and efficiency of the learning-based loop invariant inference algorithm in [14]. Experiments excerpted from Linux, SPEC2000, and Tar source codes are reported.

1 Introduction

One way to prove that an annotated loop satisfies its pre- and post-conditions is by giving loop invariants. In an annotated loop, pre- and post-conditions specify intended effects of the loop. The actual behavior of the annotated loop however does not necessarily conform to its specification. Through loop invariants, verification tools can check whether the annotated loop fulfills its specification automatically [10, 5].

Finding loop invariants is tedious and sometimes requires intelligence. Recently, an automated technique based on algorithmic learning and predicate abstraction is proposed [14]. Given a fixed set of atomic predicates and an annotated loop, the learning-based technique can infer a quantifier-free loop invariant generated by the given atomic predicates. By employing a learning algorithm and a mechanical teacher, the new technique is able to generate loop invariants without constructing abstract models nor computing fixed points. It gives a new invariant generation framework that can be less sensitive to the number of atomic predicates than traditional techniques.

As in other techniques based on predicate abstraction, the selection of atomic predicates is crucial to the effectiveness of the learning-based technique. Oftentimes, users extract atomic predicates from program texts heuristically. If this simple strategy does not yield necessary atomic predicates to express loop invariants, the loop invariant inference algorithm will not be able to infer a loop invariant. Even when the heuristic does give necessary atomic predicates, it may select too many redundant predicates and impede the efficiency of loop invariant inference algorithm.

* This work was supported by the Engineering Research Center of Excellence Program of Korea Ministry of Education, Science and Technology(MEST) / National Research Foundation of Korea(NRF) (Grant 2010-0001717), National Science Council of Taiwan Grant Numbers 97-2221-E-001-003-MY3 and 97-2221-E-001-006-MY3, the FORMES Project within LIAMA Consortium, and the French ANR project SIVES ANR-08-BLAN-0326-01.

One way to circumvent this problem is to generate atomic predicates by need. Several techniques have been developed to synthesize atomic predicates by interpolation [11, 12, 17, 18, 9]. Let A and B be logic formulae. An interpolant I of A and B is a formula such that $A \Rightarrow I$ and $I \wedge B$ is inconsistent. Moreover, the non-logical symbols in I must occur in both A and B . By Craig's interpolation theorem, an interpolant I always exists for any first-order formulae A and B when $A \wedge B$ is inconsistent [6]. The interpolant I can be seen as a concise summary of A with respect to B . Indeed, interpolants have been used to synthesize atomic predicates for predicate abstraction refinement in software model checking [11, 12, 17, 18, 9].

Inspired by the refinement technique in software model checking, we develop an interpolation-based technique to synthesize atomic predicates in the context of loop invariant inference. Our algorithm does not add new atomic predicates by interpolating invalid execution paths in control flow graphs. We instead interpolate the loop body with purported loop invariants from the learning algorithm. Our technique can improve the effectiveness and efficiency of the learning-based loop invariant inference algorithm in [14]. Constructing sets of atomic predicates can be fully automatic and on-demand.

Example. Consider the following annotated loop:

$$\{ n \geq 0 \wedge x = n \wedge y = n \} \text{ while } x > 0 \text{ do } x = x - 1; y = y - 1 \text{ done } \{ x + y = 0 \}$$

Assume that variables x and y both have the value $n \geq 0$ before entering the loop. In the loop body, each variable is decremented by one until the variable x is zero. We want to show that $x + y$ is zero after executing the loop. Note that the predicate $x = y$ is implicitly implied by the loop. The program text however does not reveal this equality explicitly. Moreover, atomic predicates from the program text can not express loop invariants that establish the specification. Using atomic predicates in the program text does not give necessary atomic predicates.

Any loop invariant must be weaker than the pre-condition and stronger than the disjunction of the loop guard and the post-condition. We use the atomic predicates in an interpolant of $n \geq 0 \wedge x = n \wedge y = n$ and $\neg(x + y = 0 \vee x > 0)$ to obtain the initial atomic predicates $\{x = y, 2y \geq 0\}$. Observe that the interpolation theorem is able to synthesize the implicit predicate $x = y$. In fact, $x = y \wedge x \geq 0$ is a loop invariant that establishes the specification of the loop.

Related Work. Loop invariant inference using algorithmic learning is introduced in [14]. In [15], the learning-based technique is extended to quantified loop invariants. Both algorithms require users to provide atomic predicates. The present work addresses this problem for the case of quantifier-free loop invariants.

Many interpolation algorithms and their implementations are available [17, 3, 7]. Interpolation-based techniques for predicate refinement in software model checking are proposed in [11, 12, 18, 9, 13]. Abstract models used in these techniques however may require excessive invocations to theorem provers. Another interpolation-based technique for first-order invariants is developed in [19]. The paramodulation-based technique does not construct abstract models. It however only generates invariants in first-order logic with equality. A template-based predicate generation technique for quantified invariants

is proposed [20]. The technique reduces the invariant inference problem to constraint programming and generates predicates in user-provided templates.

This paper is organized as follows. After Introduction, preliminaries are given in Section 2. We review the learning-based loop invariant inference framework in Section 3. Our technical results are presented in Section 4. Section 5 gives the loop invariant inference algorithm with automatic predicate generation. We report our experimental results in Section 6. Section 7 concludes this work.

2 Preliminaries

Let QF denote the quantifier-free logic with equality, linear inequality, and uninterpreted functions [17, 18]. Define the *domain* $\mathbb{D} = \mathbb{Q} \cup \mathbb{B}$ where \mathbb{Q} is the set of rational numbers and $\mathbb{B} = \{F, T\}$ is the Boolean domain. Fix a set X of variables. A *valuation* over X is a function from X to \mathbb{D} . The class of valuations over X is denoted by Val_X . For any formula $\theta \in QF$ and valuation ν over free variables in θ , θ is *satisfied* by ν (written $\nu \models \theta$) if θ evaluates to T under ν ; θ is *inconsistent* if θ is not satisfied by any valuation. Given a formula $\theta \in QF$, a *satisfiability modulo theories (SMT) solver* returns a satisfying valuation ν of θ if θ is not inconsistent [8, 16].

For $\theta \in QF$, we denote the set of non-logical symbols occurred in θ by $\sigma(\theta)$. Let $\Theta = [\theta_1, \dots, \theta_m]$ be a sequence with $\theta_i \in QF$ for $1 \leq i \leq m$. The sequence Θ is *inconsistent* if $\theta_1 \wedge \theta_2 \wedge \dots \wedge \theta_m$ is inconsistent. The sequence $\Lambda = [\lambda_0, \lambda_1, \dots, \lambda_m]$ of quantifier-free formulae is an *inductive interpolant* of Θ if

- $\lambda_0 = T$ and $\lambda_m = F$;
- for all $1 \leq i \leq m$, $\lambda_{i-1} \wedge \theta_i \Rightarrow \lambda_i$; and
- for all $1 \leq i < m$, $\sigma(\lambda_i) \subseteq \sigma(\theta_i) \cap \sigma(\theta_{i+1})$.

The interpolation theorem states that an inductive interpolant exists for any inconsistent sequence [6, 17, 18]. We consider the following imperative language in this paper:

$$\begin{aligned} \text{Stmt} &\triangleq \text{nop} \mid \text{Stmt}; \text{Stmt} \mid x := \text{Exp} \mid x := \text{nondet} \mid \text{if BExp then Stmt else Stmt} \\ \text{Exp} &\triangleq n \mid x \mid \text{Exp} + \text{Exp} \mid \text{Exp} - \text{Exp} \\ \text{BExp} &\triangleq F \mid x \mid \neg \text{BExp} \mid \text{BExp} \wedge \text{BExp} \mid \text{Exp} < \text{Exp} \mid \text{Exp} = \text{Exp} \end{aligned}$$

Two basic types are available: natural numbers and Booleans. A term in Exp is a natural number; a term in BExp is of Boolean type. The keyword *nondet* denotes an arbitrary value in the type of the assigned variable. An *annotated loop* is of the form:

$$\{\delta\} \text{ while } \kappa \text{ do } S_1; S_2; \dots; S_m \text{ done } \{\epsilon\}$$

The BExp formula κ is the *loop guard*. The BExp formulae δ and ϵ are the *precondition* and *postcondition* of the annotated loop respectively.

Define $X^{(k)} = \{x^{(k)} : x \in X\}$. For any term e over X , define $e^{(k)} = e[X \mapsto X^{(k)}]$. A *transition formula* $\llbracket S \rrbracket$ for a statement S is a first-order formula over variables $X^{(0)} \cup X^{(1)}$ defined as follows.

$$\begin{aligned}
 \llbracket \text{nop} \rrbracket &\triangleq \bigwedge_{x \in X} x^{(1)} = x^{(0)} & \llbracket x := \text{nondet} \rrbracket &\triangleq \bigwedge_{y \in X \setminus \{x\}} y^{(1)} = y^{(0)} \\
 \llbracket x := e \rrbracket &\triangleq x^{(1)} = e^{(0)} \wedge \bigwedge_{y \in X \setminus \{x\}} y^{(1)} = y^{(0)} \\
 \llbracket S_0; S_1 \rrbracket &\triangleq \exists X. \llbracket S_0 \rrbracket [X^{(1)} \mapsto X] \wedge \llbracket S_1 \rrbracket [X^{(0)} \mapsto X] \\
 \llbracket \text{if } p \text{ then } S_0 \text{ else } S_1 \rrbracket &\triangleq (p^{(0)} \wedge \llbracket S_0 \rrbracket) \vee (\neg p^{(0)} \wedge \llbracket S_1 \rrbracket)
 \end{aligned}$$

Let ν and ν' be valuations, and S a statement. We write $\nu \xrightarrow{S} \nu'$ if $\llbracket S \rrbracket$ evaluates to true by assigning $\nu(x)$ and $\nu'(x)$ to $x^{(0)}$ and $x^{(1)}$ for each $x \in X$ respectively. Given a sequence of statements $S_1; S_2; \dots; S_m$, a *program execution* $\nu_0 \xrightarrow{S_1} \nu_1 \xrightarrow{S_2} \dots \xrightarrow{S_m} \nu_m$ is a sequence $[\nu_0, \nu_1, \dots, \nu_m]$ of valuations such that $\nu_i \xrightarrow{S_i} \nu_{i+1}$ for $0 \leq i < m$.

A *precondition* $Pre(\theta : S_1; S_2; \dots; S_m)$ for $\theta \in QF$ with respect to the statement $S_1; S_2; \dots; S_m$ is a first-order formula that entails θ after executing the statement $S_1; S_2; \dots; S_m$. Given an annotated loop $\{\delta\}$ while κ do $S_1; S_2; \dots; S_m$ done $\{\epsilon\}$, the *loop invariant inference problem* is to compute a formula $\iota \in QF$ satisfying (1) $\delta \Rightarrow \iota$; (2) $\iota \wedge \neg \kappa \Rightarrow \epsilon$; and (3) $\iota \wedge \kappa \Rightarrow Pre(\iota : S_1; S_2; \dots; S_m)$. Observe that the condition (2) is equivalent to $\iota \Rightarrow \epsilon \vee \kappa$. The first two conditions specify necessary and sufficient conditions of any loop invariants respectively. The formulae δ and $\epsilon \vee \kappa$ are called the *strongest* and *weakest approximations* to loop invariants respectively.

3 Inferring Loop Invariants with Algorithmic Learning

Given an annotated loop $\{\delta\}$ while κ do $S_1; S_2; \dots; S_m$ done $\{\epsilon\}$, we would like to infer a loop invariant to establish the pre- and post-conditions. Given a set P of atomic predicates, the work in [14] shows how to apply a learning algorithm for Boolean formulae to infer quantifier-free loop invariants freely generated by P . The authors first adopt predicate abstraction to relate quantifier-free and Boolean formulae. They then design a mechanical teacher to guide the learning algorithm to a Boolean formula whose concretization is a loop invariant.




Fig. 1. Relating QF and $Bool[B_P]$

Let $QF[P]$ denote the set of quantifier-free formulae generated from the set of atomic predicates P . Consider the set of Boolean formulae $Bool[B_P]$ generated by the set of Boolean variables $B_P \triangleq \{b_p : p \in P\}$. An *abstract valuation* is a function from B_P to \mathbb{B} . We write Val_{B_P} for the set of abstract valuations. A Boolean formula in $Bool[B_P]$ is a *canonical monomial* if it is a conjunction of literals, where each Boolean variable in B_P occurs exactly once. Formulae in $QF[P]$ and $Bool[B_P]$ are related by the following functions [14] (Figure 1):

$$\begin{aligned}
\gamma(\beta) &\triangleq \beta[B_P \mapsto P] \\
\alpha(\theta) &\triangleq \bigvee \{ \beta \in \text{Bool}[B_P] : \beta \text{ is a canonical monomial and } \theta \wedge \gamma(\beta) \text{ is satisfiable} \} \\
\gamma^*(\mu) &\triangleq \bigwedge_{\mu(b_p)=T} \{p\} \wedge \bigwedge_{\mu(b_p)=F} \{\neg p\} \\
\alpha^*(\nu) &\triangleq \mu \text{ where } \mu(b_p) = \begin{cases} T & \text{if } \nu \models p \\ F & \text{if } \nu \not\models p \end{cases}
\end{aligned}$$

Consider, for instance, $P = \{n \geq 0, x = n, y = n\}$ and $B_P = \{b_{n \geq 0}, b_{x=n}, b_{y=n}\}$. We have $\gamma(b_{n \geq 0} \wedge \neg b_{x=n}) = n \geq 0 \wedge \neg(x = n)$ and

$$\begin{aligned}
\alpha(\neg(x = y)) &= (b_{n \geq 0} \wedge b_{x=n} \wedge \neg b_{y=n}) \vee (b_{n \geq 0} \wedge \neg b_{x=n} \wedge b_{y=n}) \vee \\
&\quad (b_{n \geq 0} \wedge \neg b_{x=n} \wedge \neg b_{y=n}) \vee (\neg b_{n \geq 0} \wedge b_{x=n} \wedge \neg b_{y=n}) \vee \\
&\quad (\neg b_{n \geq 0} \wedge \neg b_{x=n} \wedge b_{y=n}) \vee (\neg b_{n \geq 0} \wedge \neg b_{x=n} \wedge \neg b_{y=n}).
\end{aligned}$$

Moreover, $\alpha^*(\nu)(b_{n \geq 0}) = \alpha^*(\nu)(b_{x=n}) = \alpha^*(\nu)(b_{y=n}) = T$ when $\nu(n) = \nu(x) = \nu(y) = 1$. And $\gamma^*(\mu) = n \geq 0 \wedge x = n \wedge \neg(y = n)$ when $\mu(b_{n \geq 0}) = \mu(b_{x=n}) = T$ but $\mu(b_{y=n}) = F$. Observe that the pair (α, γ) forms the Galois correspondence in Cartesian predicate abstraction [2].

After formulae in QF and valuations in Val_X are abstracted to $\text{Bool}[B_P]$ and Val_{B_P} respectively, a learning algorithm is used to infer abstractions of loop invariants. Let ξ be an unknown *target* Boolean formula in $\text{Bool}[B_P]$. A *learning algorithm* computes a representation of the target ξ by interacting with a teacher. The *teacher* should answer the following queries [1, 4]:

- *Membership queries.* Let $\mu \in Val_{B_P}$ be an abstract valuation. The membership query $MEM(\mu)$ asks if the unknown target ξ is satisfied by μ . If so, the teacher answers *YES*; otherwise, *NO*.
- *Equivalence queries.* Let $\beta \in \text{Bool}[B_P]$ be an *abstract conjecture*. The equivalence query $EQ(\beta)$ asks if β is equivalent to the unknown target ξ . If so, the teacher answers *YES*. Otherwise, the teacher gives an abstract valuation μ such that the exclusive disjunction of β and ξ is satisfied by μ . The abstract valuation μ is called an *abstract counterexample*.

With predicate abstraction and a learning algorithm for Boolean formulae at hand, it remains to design a mechanical teacher to guide the learning algorithm to the abstraction of a loop invariant. The key idea in [14] is to exploit approximations to loop invariants. An *under-approximation* to loop invariants is a quantifier-free formula $\underline{\iota}$ which is stronger than some loop invariants of the given annotated loop; an *over-approximation* is a quantifier-free formula $\bar{\iota}$ which is weaker than some loop invariants.

To see how approximations to loop invariants can be used in the design of the mechanical teacher, let us consider an equivalence query $EQ(\beta)$. On the abstract conjecture $\beta \in \text{Bool}[B_P]$, the mechanical teacher computes the corresponding quantifier-free formula $\theta = \gamma(\beta)$. It then checks if θ is a loop invariant. If so, we are done. Otherwise, the algorithm compares θ with approximations to loop invariants. If θ is stronger than the under-approximation or weaker than the over-approximation, a valuation ν satisfying $\neg(\underline{\iota} \Rightarrow \theta)$ or $\neg(\theta \Rightarrow \bar{\iota})$ can be obtained from an SMT solver. The abstract




Fig. 2. Learning-based Framework

valuation $\alpha^*(\nu)$ gives an abstract counterexample. Approximations to loop invariants can also be used to answer membership queries. For a membership query $MEM(\mu)$ with $\mu \in Val_{BP}$, the mechanical teacher computes its concretization $\theta = \gamma^*(\mu)$. It returns *YES* if $\theta \Rightarrow \underline{L}$; it returns *NO* if $\theta \not\Rightarrow \bar{L}$. Otherwise, a random answer is returned.

Figure 2 shows the learning-based loop invariant inference framework. In the framework, a learning algorithm is used to drive the search of loop invariants. It “learns” an unknown loop invariant by inquiring a mechanical teacher. The mechanical teacher of course does not know any loop invariant. It nevertheless can try to answer these queries by the information derived from program texts. In this case, approximations to loop invariants are used. Observe the simplicity of the learning-based framework. By employing a learning algorithm, it suffices to design a mechanical teacher to find loop invariants. Moreover, the new framework does not construct abstract models nor compute fixed points. It can be more scalable than traditional techniques.

4 Predicate Generation by Interpolation

One drawback in the learning-based approach to loop invariant inference is to require a set of atomic predicates. It is essential that at least one quantifier-free loop invariant is representable by the given set P of atomic predicates. Otherwise, concretization of formulae in $Bool[B_P]$ cannot be loop invariants. The mechanical teacher never answers *YES* to equivalence queries. To address this problem, we will synthesize new atomic predicates for the learning-based loop invariant inference framework progressively.

The interpolation theorem is essential to our predicate generation technique [6, 19, 18, 12]. Let $\Theta = [\theta_1, \theta_2, \dots, \theta_m]$ be an inconsistent sequence of quantifier-free formula and $\Lambda = [\lambda_0, \lambda_1, \lambda_2, \dots, \lambda_m]$ its inductive interpolant. By definition, $\theta_1 \Rightarrow \lambda_1$. Assume $\theta_1 \wedge \theta_2 \wedge \dots \wedge \theta_i \Rightarrow \lambda_i$. We have $\theta_1 \wedge \theta_2 \wedge \dots \wedge \theta_{i+1} \Rightarrow \lambda_{i+1}$ since $\lambda_i \wedge \theta_{i+1} \Rightarrow \lambda_{i+1}$. Thus, λ_i is an over-approximation to $\theta_1 \wedge \theta_2 \wedge \dots \wedge \theta_i$ for $0 \leq i \leq m$. Moreover, $\sigma(\lambda_i) \subseteq \sigma(\theta_i) \cap \sigma(\theta_{i+1})$. Hence λ_i can be seen as a concise summary of $\theta_1 \wedge \theta_2 \wedge \dots \wedge \theta_i$ with restricted symbols. Since each λ_i is written in a less expressive vocabulary, new atomic predicates among variables can be synthesized. We therefore apply the interpolation theorem to synthesize new atomic predicates and refine the abstraction.

Our predicate generation technique consists of three components. Before the learning algorithm is invoked, an initial set of atomic predicates is computed (Section 4.1). When

the learning algorithm is failing to infer loop invariants, new atomic predicates are generated to refine the abstraction (Section 4.2). Lastly, conflicting answers to queries may incur from predicate abstraction. We further refine the abstraction with these conflicting answers (Section 4.3). Throughout this section, we consider the annotated loop $\{\delta\}$ while κ do $S_1; S_2; \dots; S_m$ done $\{\epsilon\}$ with the under-approximation $\underline{\iota}$ and over-approximation $\bar{\iota}$.

4.1 Initial Atomic Predicates

The under- and over-approximations to loop invariants must satisfy $\underline{\iota} \Rightarrow \bar{\iota}$. Otherwise, there cannot be any loop invariant ι such that $\underline{\iota} \Rightarrow \iota$ and $\iota \Rightarrow \bar{\iota}$. Thus, the sequence $[\underline{\iota}, \bar{\iota}]$ is inconsistent. For any interpolant $[T, \lambda, F]$ of $[\underline{\iota}, \bar{\iota}]$, we have $\underline{\iota} \Rightarrow \lambda$ and $\lambda \Rightarrow \bar{\iota}$. The quantifier-free formula λ can be a loop invariant if it satisfies $\lambda \wedge \kappa \Rightarrow \text{Pre}(\lambda : S_1; S_2; \dots; S_m)$. It is however unlikely that λ happens to be a loop invariant. Yet our loop invariant inference algorithm can generalize λ by taking the atomic predicates in λ as the initial atomic predicates. The learning algorithm will try to infer a loop invariant freely generated by these atomic predicates.

4.2 Atomic Predicates from Incorrect Conjectures

Consider an equivalence query $EQ(\beta)$ where $\beta \in \text{Bool}[B_P]$ is an abstract conjecture. If the concretization $\theta = \gamma(\beta)$ is not a loop invariant, we interpolate the loop body with the incorrect conjecture θ . For any quantifier-free formula θ over variables $X^{(0)} \cup X^{(1)}$, define $\theta^{(k)} = \theta[X^{(0)} \mapsto X^{(k)}, X^{(1)} \mapsto X^{(k+1)}]$. The *desuperscripted* form of a quantifier-free formula λ over variables $X^{(k)}$ is $\lambda[X^{(k)} \mapsto X]$. Moreover, if ν is a valuation over $X^{(0)} \cup \dots \cup X^{(m)}$, $\nu \downarrow_{X^{(k)}}$ represents a valuation over X such that $\nu \downarrow_{X^{(k)}}(x) = \nu(x^{(k)})$ for $x \in X$. Let ϕ and ψ be quantifier-free formulae over X . Define the following sequence:

$$\Xi(\phi, S_1, \dots, S_m, \psi) \triangleq [\phi^{(0)}, \llbracket S_1 \rrbracket^{(0)}, \llbracket S_2 \rrbracket^{(1)}, \dots, \llbracket S_m \rrbracket^{(m-1)}, \neg\psi^{(m)}].$$

Observe that

- $\phi^{(0)}$ and $\llbracket S_1 \rrbracket^{(0)}$ share the variables $X^{(0)}$;
- $\llbracket S_m \rrbracket^{(m-1)}$ and $\neg\psi^{(m)}$ share the variables $X^{(m)}$; and
- $\llbracket S_i \rrbracket^{(i-1)}$ and $\llbracket S_{i+1} \rrbracket^{(i)}$ share the variables $X^{(i)}$ for $1 \leq i < m$.

Starting from the program states satisfying $\phi^{(0)}$, the formula

$$\phi^{(0)} \wedge \llbracket S_1 \rrbracket^{(0)} \wedge \llbracket S_2 \rrbracket^{(1)} \wedge \dots \wedge \llbracket S_i \rrbracket^{(i-1)}$$

characterizes the images of $\phi^{(0)}$ during the execution of $S_1; S_2; \dots; S_i$.

Lemma 1. *Let X denote the set of variables in the statement $S_1; S_2; \dots; S_i$, and ϕ a quantifier-free formula over X . For any valuation ν over $X^{(0)} \cup X^{(1)} \cup \dots \cup X^{(i)}$, the formula $\phi^{(0)} \wedge \llbracket S_1 \rrbracket^{(0)} \wedge \llbracket S_2 \rrbracket^{(1)} \wedge \dots \wedge \llbracket S_i \rrbracket^{(i-1)}$ is satisfied by ν if and only if $\nu \downarrow_{X^{(0)}} \xrightarrow{S_1} \nu \downarrow_{X^{(1)}} \xrightarrow{S_2} \dots \xrightarrow{S_i} \nu \downarrow_{X^{(i-1)}}$ is a program execution and $\nu \downarrow_{X^{(0)}} \models \phi$.*

By definition, $\phi \Rightarrow \text{Pre}(\psi : S_1; S_2; \dots; S_m)$ implies that the image of ϕ must satisfy ψ after the execution of $S_1; S_2; \dots; S_m$. The sequence $\Xi(\phi, S_1, \dots, S_m, \psi)$ is inconsistent if $\phi \Rightarrow \text{Pre}(\psi : S_1; S_2; \dots; S_m)$. The following proposition will be handy.

Proposition 1. *Let $S_1; S_2; \dots; S_m$ be a sequence of statements. For any ϕ with $\phi \Rightarrow \text{Pre}(\psi : S_1; S_2; \dots; S_m)$, $\Xi(\phi, S_1, \dots, S_m, \psi)$ has an inductive interpolant¹.*

Let $A = [T, \lambda_1, \lambda_2, \dots, \lambda_{m+1}, F]$ be an inductive interpolant of $\Xi(\phi, S_1, \dots, S_m, \psi)$. Recall that λ_i is a quantifier-free formula over $X^{(i-1)}$ for $1 \leq i \leq m+1$. It is also an over-approximation to the image of ϕ after executing $S_1; S_2; \dots; S_{i-1}$. Proposition 1 can be used to generate new atomic predicates. One simply finds a pair of quantifier-free formulae ϕ and ψ with $\phi \Rightarrow \text{Pre}(\psi : S_1; S_2; \dots; S_m)$, applies the interpolation theorem, and collects desuperscripted atomic predicates in an inductive interpolant of $\Xi(\phi, S_1, \dots, S_m, \psi)$. In the following, we show how to obtain such pairs with under and over-approximations to loop invariants.

Interpolating Over-Approximation. It is not hard to see that an over-approximation to loop invariants characterizes loop invariants after the execution of the loop body. Recall that $\iota \Rightarrow \bar{\tau}$ for some loop invariant ι . Moreover, $\iota \wedge \kappa \Rightarrow \text{Pre}(\iota : S_1; S_2; \dots; S_m)$. By the monotonicity of $\text{Pre}(\bullet : S_1; S_2; \dots; S_m)$, we have $\iota \wedge \kappa \Rightarrow \text{Pre}(\bar{\tau} : S_1; S_2; \dots; S_m)$.

Proposition 2. *Let $\bar{\tau}$ be an over-approximation to loop invariants of the annotated loop $\{\delta\}$ while κ do $S_1; S_2; \dots; S_m$ done $\{\epsilon\}$. For any loop invariant ι with $\iota \Rightarrow \bar{\tau}$, $\iota \wedge \kappa \Rightarrow \text{Pre}(\bar{\tau} : S_1; S_2; \dots; S_m)$.*

Proposition 2 gives a necessary condition to loop invariants of interest. Recall that $\theta = \gamma(\beta)$ is an incorrect conjecture of loop invariants. If $\nu \models \neg(\theta \wedge \kappa \Rightarrow \text{Pre}(\bar{\tau} : S_1; S_2; \dots; S_m))$, the mechanical teacher returns the abstract counterexample $\alpha^*(\nu)$. Otherwise, Proposition 1 is applicable with the pair $\theta \wedge \kappa$ and $\bar{\tau}$.

Corollary 1. *Let $\bar{\tau}$ be an over-approximation to loop invariants of the annotated loop $\{\delta\}$ while κ do $S_1; S_2; \dots; S_m$ done $\{\epsilon\}$. For any θ with $\theta \wedge \kappa \Rightarrow \text{Pre}(\bar{\tau} : S_1; S_2; \dots; S_m)$, the sequence $\Xi(\theta \wedge \kappa, S_1, S_2, \dots, S_m, \bar{\tau})$ has an inductive interpolant.*

Interpolating Under-Approximation. For under-approximations, there is no necessary condition. Nevertheless, Proposition 1 is applicable with the pair $\underline{\iota} \wedge \kappa$ and θ .

Corollary 2. *Let $\underline{\iota}$ be an under-approximation to loop invariants of the annotated loop $\{\delta\}$ while κ do $S_1; S_2; \dots; S_m$ done $\{\epsilon\}$. For any θ with $\underline{\iota} \wedge \kappa \Rightarrow \text{Pre}(\theta : S_1; S_2; \dots; S_m)$, the sequence $\Xi(\underline{\iota} \wedge \kappa, S_1, S_2, \dots, S_m, \theta)$ has an inductive interpolant.*

Generating atomic predicates from an incorrect conjecture θ should now be clear (Algorithm 1). Assuming that the incorrect conjecture satisfies the necessary condition in Proposition 2, we simply collect all desuperscripted atomic predicates in an inductive interpolant of $\Xi(\theta \wedge \kappa, S_1, S_2, \dots, S_m, \bar{\tau})$ (Corollary 1). More atomic predicates can be obtained from an inductive interpolant of $\Xi(\underline{\iota} \wedge \kappa, S_1, S_2, \dots, S_m, \theta)$ if additionally $\underline{\iota} \wedge \kappa \Rightarrow \text{Pre}(\theta : S_1; S_2; \dots; S_m)$ (Corollary 2).

¹ The existential quantifiers in $\llbracket S; S' \rrbracket$ are eliminated by introducing fresh variables.


```

/* {δ} while κ do S1;⋯;Sm done {ε} : an annotated loop */
/*  $\underline{L}, \bar{\tau}$  : under- and over-approximations to loop invariants */
Input: a formula  $\theta \in QF[P]$  such that  $\theta \wedge \kappa \Rightarrow Pre(\bar{\tau} : S_1; S_2; \dots; S_m)$ 
Output: a set of atomic predicates
I := an inductive interpolant of  $\Xi(\theta \wedge \kappa, S_1, S_2, \dots, S_m, \bar{\tau})$ ;
Q := desuperscripted atomic predicates in I;
if  $\underline{L} \wedge \kappa \Rightarrow Pre(\theta : S_1; S_2; \dots; S_m)$  then
  J := an inductive interpolants of  $\Xi(\underline{L} \wedge \kappa, S_1, S_2, \dots, S_m, \theta)$ ;
  R := desuperscripted atomic predicates in J;
  Q := Q  $\cup$  R;
end
return Q

```

Algorithm 1. PredicatesFromConjecture (θ)

4.3 Atomic Predicates from Conflicting Abstract Counterexamples

Because of the abstraction, conflicting abstract counterexamples may be given to the learning algorithm. Consider the example in Section 1. Recall that $n \geq 0 \wedge x = n \wedge y = n$ and $x + y = 0 \vee x > 0$ are the under- and over-approximations respectively. Suppose there is only one atomic predicate $y = 0$. The learning algorithm tries to infer a Boolean formula $\lambda \in Bool[b_{y=0}]$. Let us resolve the equivalence queries $EQ(T)$ and $EQ(F)$. On the equivalence query $EQ(F)$, we check if F is weaker than the under-approximation by an SMT solver. It is not, and the SMT solver gives the valuation $\nu_0(n) = \nu_0(x) = \nu_0(y) = 1$ as a witness. Applying the abstraction function α^* to ν_0 , the mechanical teacher returns the abstract counterexample $b_{y=0} \mapsto F$. The abstract counterexample is intended to notify that the target formula λ and F have different truth values when $b_{y=0}$ is F . That is, λ is satisfied by the valuation $b_{y=0} \mapsto F$.

On the equivalence query $EQ(T)$, the mechanical teacher checks if T is stronger than the over-approximation. It is not, and the SMT solver now returns the valuation $\nu_1(x) = 0, \nu_1(y) = 1$ as a witness. The mechanical teacher in turn computes $b_{y=0} \mapsto F$ as the corresponding abstract counterexample. The abstract counterexample notifies that the target formula λ and T have different truth values when $b_{y=0}$ is F . That is, λ is not satisfied by the valuation $b_{y=0} \mapsto F$. Yet the target formula λ cannot be satisfied and unsatisfied by the valuation $b_{y=0} \mapsto F$. We have conflicting abstract counterexamples.

Such conflicting abstract counterexamples arise because the abstraction is too coarse. This gives us another chance to refine the abstraction. Define

$$\Gamma(\nu) \triangleq \bigwedge_{x \in X} x = \nu(x).$$

The function $\Gamma(\nu)$ specifies the valuation ν in QF (Figure 1). For distinct valuations ν and ν' , $\Gamma(\nu) \wedge \Gamma(\nu')$ is inconsistent. For instance, $\Gamma(\nu_0) = (n = 1) \wedge (x = 1) \wedge (y = 1)$, $\Gamma(\nu_1) = (x = 0) \wedge (y = 1)$, and $\Gamma(\nu_1) \wedge \Gamma(\nu_0)$ is inconsistent.

Algorithm 2 generates atomic predicates from conflicting abstract counterexamples. Let ν and ν' be distinct valuations in Val_X . We compute formulae $\chi = \Gamma(\nu)$ and $\chi' = \Gamma(\nu')$. Since ν and ν' are conflicting, they correspond to the same abstract valuation $\alpha^*(\nu) = \alpha^*(\nu')$. Let $\rho = \gamma^*(\alpha^*(\nu))$. We have $\chi \Rightarrow \rho$ and $\chi' \Rightarrow \rho$ [14]. Recall that

Input: distinct valuations ν and ν' such that $\alpha^*(\nu) = \alpha^*(\nu')$

Output: a set of atomic predicates

$\chi := \Gamma(\nu)$;

$\chi' := \Gamma(\nu')$;

/ $\chi \wedge \chi'$ is inconsistent*

**/*

$\rho := \gamma^*(\alpha^*(\nu))$;

$Q :=$ atomic predicates in an inductive interpolant of $[\chi, \chi' \vee \neg\rho]$;

return Q ;

Algorithm 2. PredicatesFromConflict (ν, ν')

$\chi \wedge \chi'$ is inconsistent. $[\chi, \chi' \vee \neg\rho]$ is also inconsistent for $\chi \Rightarrow \rho$. Algorithm 2 returns atomic predicates in an inductive interpolant of $[\chi, \chi' \vee \neg\rho]$.

5 Algorithm

Our loop invariant inference algorithm is given in Algorithm 3. For an annotated loop $\{\delta\}$ **while** κ **do** $S_1; S_2; \dots; S_m$ **done** $\{\epsilon\}$, we heuristically choose $\delta \vee \epsilon$ and $\epsilon \vee \kappa$ as the under- and over-approximations respectively. Note that the under-approximation is different from the strongest approximation δ . It is reported that the approximations $\delta \vee \epsilon$ and $\epsilon \vee \kappa$ are more effective in resolving queries [14].

/ $\{\delta\}$ **while** κ **do** $S_1; S_2; \dots; S_m$ **done** $\{\epsilon\}$: an annotated loop*

Output: a loop invariant for the annotated loop

$\underline{\delta} := \delta \vee \epsilon$;

$\bar{\epsilon} := \epsilon \vee \kappa$;

$P :=$ atomic predicates in an inductive interpolant of $[\underline{\delta}, \bar{\epsilon}]$;

repeat

try

call a learning algorithm for Boolean formulae where membership and equivalence queries are resolved by Algorithms 4 and 5 respectively;

catch conflict abstract counterexamples \rightarrow

find distinct valuations ν and ν' such that $\alpha^*(\nu) = \alpha^*(\nu')$;

$P := P \cup \text{PredicatesFromConflict}(\nu, \nu')$;

until a loop invariant is found ;

Algorithm 3. Loop Invariant Inference

We compute the initial atomic predicates by interpolating $\underline{\delta}$ and $\bar{\epsilon}$ (Section 4.1). The main loop invokes a learning algorithm. It resolves membership and equivalence queries from the learning algorithm by under- and over-approximations (detailed later). If there is a conflict, the loop invariant inference algorithm adds more atomic predicates by Algorithm 2. Then the main loop reiterates with the new set of atomic predicates.

For membership queries, we compare the concretization of the abstract valuation with approximations to loop invariants (Algorithm 4). The mechanical teacher returns *NO* when the concretization is inconsistent. If the concretization is stronger than the under-approximation, the mechanical teacher returns *YES*; if the concretization is

```

/*  $\underline{\iota}, \bar{\iota}$  : under- and over-approximations to loop invariants */
Input: a membership query  $MEM(\mu)$  with  $\mu \in Val_{B_P}$ 
Output: YES or NO
 $\theta := \gamma^*(\mu)$ ;
if  $\theta$  is inconsistent then return NO;
if  $\theta \Rightarrow \underline{\iota}$  then return YES;
if  $\nu \models \neg(\theta \Rightarrow \bar{\iota})$  then return NO;
return YES or NO randomly;

```

Algorithm 4. Membership Query Resolution

```

/*  $\tau$  : a threshold to generate new atomic predicates */
/*  $\{\delta\}$  while  $\kappa$  do  $S_1; S_2; \dots; S_m$  done  $\{\epsilon\}$  : an annotated loop */
/*  $\underline{\iota}, \bar{\iota}$  : under- and over-approximations to loop invariants */
Input: an equivalence query  $EQ(\beta)$  with  $\beta \in Bool[B_P]$ 
Output: YES or an abstract counterexample
 $\theta := \gamma(\beta)$ ;
if  $\delta \Rightarrow \theta$  and  $\theta \Rightarrow \epsilon \vee \kappa$  and  $\theta \wedge \kappa \Rightarrow Pre(\theta : S_1; S_2; \dots; S_m)$  then return YES;
if  $\nu \models \neg(\underline{\iota} \Rightarrow \theta)$  or  $\nu \models \neg(\theta \Rightarrow \bar{\iota})$  or  $\nu \models \neg(\theta \wedge \kappa \Rightarrow Pre(\bar{\iota} : S_1; S_2; \dots; S_m))$  then
  record  $\nu$ ; return  $\alpha^*(\nu)$ ;
if the number of random abstract counterexamples  $\leq \tau$  then
  return a random abstract counterexample;
else
   $P := P \cup PredicatesFromConjecture(\theta)$ ;  $\tau := \lceil 1.3^{|P|} \rceil$ ; reiterate the main loop;
end

```

Algorithm 5. Equivalence Query Resolution

weaker than the over-approximation, it returns *NO*. Otherwise, a random answer is returned [14].

The equivalence query resolution algorithm is given in Algorithm 5. For any equivalence query, the mechanical teacher checks if the concretization of the abstract conjecture is a loop invariant. If so, it returns *YES* and concludes the loop invariant inference algorithm. Otherwise, the mechanical teacher compares the concretization of the abstract conjecture with approximations to loop invariants. If the concretization is stronger than the under-approximation, weaker than the over-approximation, or it does not satisfy the necessary condition given in Proposition 2, an abstract counterexample is returned after recording the witness valuation [14, 15]. The witnessing valuations are needed to synthesize atomic predicates when conflicts occur.

If the concretization is not a loop invariant and falls between both approximations to loop invariants, there are two possibilities. The current set of atomic predicates is sufficient to express a loop invariant; the learning algorithm just needs a few more iterations to infer a solution. Or, the current atomic predicates are insufficient to express any loop invariant; the learning algorithm cannot derive a solution with these predicates. Since we cannot tell which scenario arises, a threshold is deployed heuristically. If the number of random abstract counterexamples is less than the threshold, we give the learning algorithm more time to find a loop invariant. Only when the number of random abstract counterexamples exceeds the threshold, can we synthesize more atomic predicates for abstraction refinement. Intuitively, the current atomic predicates are likely

to be insufficient if lots of random abstract counterexamples have been generated. In this case, we invoke Algorithm 2 to synthesize more atomic predicates from the incorrect conjecture, update the threshold to $\lceil 1.3^{|P|} \rceil$, and then restart the main loop.

6 Experimental Results

We have implemented the proposed technique in OCaml². In our implementation, the SMT solver YICES and the interpolating theorem prover CSISAT [3] are used for query resolution and interpolation respectively. In addition to the examples in [14], we add two more examples: `riva` is the largest loop expressible in our simple language from Linux³, and `tar` is extracted from Tar⁴. All examples are translated into annotated loops manually. Data are the average of 100 runs and collected on a 2.4GHz Intel Core2 Quad CPU with 8GB memory running Linux 2.6.31 (Table 1).

Table 1. Experimental Results.

P : # of atomic predicates, MEM : # of membership queries, EQ : # of equivalence queries, RE : # of the learning algorithm restarts, T : total elapsed time (s).

case	SIZE	PREVIOUS [14]					CURRENT					BLAST [18]	
		P	MEM	EQ	RE	T	P	MEM	EQ	RE	T	P	T
ide-ide-tape	16	6	13	7	1	0.05	4	6	5	1	0.05	21	1.31(1.07)
ide-wait-ireason	9	5	790	445	33	1.51	5	122	91	7	1.09	9	0.19(0.14)
parser	37	17	4,223	616	13	13.45	9	86	32	1	0.46	8	0.74(0.49)
riva	82	20	59	11	2	0.51	7	14	5	1	0.37	12	1.50(1.17)
tar	7	6	∞	∞	∞	∞	2	2	5	1	0.02	10	0.20(0.17)
usb-message	18	10	21	7	1	0.10	3	7	6	1	0.04	4	0.18(0.14)
vpr	8	5	16	9	2	0.05	1	1	3	1	0.01	4	0.13(0.10)

In the table, the column PREVIOUS represents the work in [14] where atomic predicates are chosen heuristically. Specifically, all atomic predicates in pre- and post-conditions, loop guards, and conditions of `if` statements are selected. The column CURRENT gives the results for our automatic predicate generation technique. Interestingly, heuristically chosen atomic predicates suffice to infer loop invariants for all examples except `tar`. For the `tar` example, the learning-based loop invariant inference algorithm fails to find a loop invariant due to ill-chosen atomic predicates. In contrast, our new algorithm is able to infer a loop invariant for the `tar` example in 0.02s. The number of atomic predicates can be significantly reduced as well. Thanks to a smaller number of atomic predicates, loop invariant inference becomes more economical in these examples. Without predicate generation, four of the six examples take more than one second. Only one of these examples takes more than one second using the new technique. Particularly, the `parser` example is improved in orders of magnitude.

² Available at <http://ropas.snu.ac.kr/tacas11/ap-gen.tar.gz>

³ In Linux 2.6.30 `drivers/video/riva/riva_hw.c:nv10CalcArbitration()`

⁴ In Tar 1.13 `src/mangle.c:extract_mangle()`

The column BLAST gives the results of lazy abstraction technique with interpolants implemented in BLAST [18]. In addition to the total elapsed time, we also show the preprocessing time in parentheses. Since the learning-based framework does not construct abstract models, our new technique outperforms BLAST in all cases but one (*ide-wait-ireason*). If we disregard the time for preprocessing in BLAST, the learning-based technique still wins three cases (*ide-ide-tape*, *tar*, *vpr*) and ties one (*usb-message*). Also note that the number of atomic predicates generated by the new technique is always smaller except *parser*. Given the simplicity of the learning-based framework, our preliminary experimental results suggest a promising outlook for further optimizations.

6.1 tar from Tar

This simple fragment is excerpted from the code for copying two buffers. M items in the source buffer are copied to the target buffer that already has N items. The variable *size* keeps the number of remaining items in the source buffer and *copy* denotes the number of items in the target buffer after the last copy. In each iteration, an arbitrary number of items are copied and the values of *size* and *copy* are updated accordingly.

```

{ size = M ∧ copy = N }
1 while size > 0 do
2   available := nondet;
3   if available > size then
4     copy := copy + available;
5     size := size - available;
6 done
{ size = 0 ⇒ copy = M + N }

```

Fig. 3. A Sample Loop in Tar

Observe that the atomic predicates in the program text cannot express any loop invariant that proves the specification. However, our new algorithm successfully finds the following loop invariant in this example:

$$M + N \leq copy + size \wedge copy + size \leq M + N$$

The loop invariant asserts that the number of items in both buffers is equal to $M + N$. It requires atomic predicates unavailable from the program text. Predicate generation is essential to find loop invariants for such tricky loops.

6.2 parser from SPEC2000 Benchmarks

For the *parser* example (Figure 4), 9 atomic predicates are generated. These atomic predicates are a subset of the 17 atomic predicates from the program text. Every loop invariant found by the loop invariant inference algorithm contains all 9 atomic predicates. This suggests that there are no redundant predicates. Few atomic predicates make loop invariants easier to comprehend. For instance, the following loop invariant summarizes the condition when *success* or *give_up* is true:

$$\begin{aligned}
& (success \vee give_up) \Rightarrow \\
& \quad (valid \neq 0 \vee cutoff = maxcost \vee words < count) \wedge \\
& \quad (\neg search \vee valid \neq 0 \vee words < count) \wedge \\
& \quad (linkages = canonical \wedge linkages \geq valid \wedge linkages \leq 5000)
\end{aligned}$$

```

{ phase = F ∧ success = F ∧ give_up = F ∧ cutoff = 0 ∧ count = 0 }
1 while ¬(success ∨ give_up) do
2   entered_phase := F;
3   if ¬phase then
4     if cutoff = 0 then cutoff := 1;
5     else if cutoff = 1 ∧ maxcost > 1 then cutoff := maxcost;
6         else phase := T; entered_phase := T; cutoff := 1000;
7     if cutoff = maxcost ∧ ¬search then give_up := T;
8   else
9     count := count + 1;
10    if count > words then give_up := T;
11    if entered_phase then count := 1;
12    linkages := nondet;
13    if linkages > 5000 then linkages := 5000;
14    canonical := 0; valid := 0;
15    if linkages ≠ 0 then
16      valid := nondet;
17      assume 0 ≤ valid ∧ valid ≤ linkages;
18      canonical := linkages;
19    if valid > 0 then success := T;
20 done
{ (valid > 0 ∨ count > words ∨ (cutoff = maxcost ∧ ¬search)) ∧
  valid ≤ linkages ∧ canonical = linkages ∧ linkages ≤ 5000 }

```

Fig. 4. A Sample Loop in SPEC2000 Benchmark PARSER

Fewer atomic predicates also lead to a smaller standard deviation of the execution time. The execution time now ranges from 0.36s to 0.58s with the standard deviation equal to 0.06. In contrast, the execution time for [14] ranges from 1.20s to 80.20s with the standard deviation equal to 14.09. By Chebyshev’s inequality, the new algorithm infers a loop invariant in one second with probability greater than 0.988. With a compact set of atomic predicates, loop invariant inference algorithm performs rather predictably.

7 Conclusions

A predicate generation technique for learning-based loop invariant inference was presented. The technique applies the interpolation theorem to synthesize atomic predicates implicitly implied by program texts. To compare the efficiency of the new technique, examples excerpted from Linux, SPEC2000, and Tar source codes were reported. The learning-based loop invariant inference algorithm is more effective and performs much better in these realistic examples.

More experiments are always needed. Especially, we would like to have more realistic examples which require implicit predicates unavailable in program texts. Additionally, loops manipulating arrays often require quantified loop invariants with linear inequalities. Extension to quantified loop invariants is also important.

Acknowledgment. The authors would like to thank Wontae Choi, Soonho Kong, and anonymous referees for their comments in improving this work.

References

1. Angluin, D.: Learning regular sets from queries and counterexamples. *Information and Computation* 75(2), 87–106 (1987)
2. Ball, T., Podolski, A., Rajamani, S.K.: Boolean and cartesian abstraction for model checking c programs. In: Margaria, T., Yi, W. (eds.) TACAS 2001. LNCS, vol. 2031, pp. 268–283. Springer, Heidelberg (2001)
3. Beyer, D., Zufferey, D., Majumdar, R.: CSISAT: Interpolation for LA+EUf. In: Gupta, A., Malik, S. (eds.) CAV 2008. LNCS, vol. 5123, pp. 304–308. Springer, Heidelberg (2008)
4. Bshouty, N.H.: Exact learning boolean functions via the monotone theory. *Information and Computation* 123, 146–153 (1995)
5. Canet, G., Cuoq, P., Monate, B.: A value analysis for c programs. In: *Source Code Analysis and Manipulation*, pp. 123–124. IEEE, Los Alamitos (2009)
6. Craig, W.: Linear reasoning, a new form of the herbrand-gentzen theorem. *J. Symb. Log.* 22(3), 250–268 (1957)
7. D’Silva, V., Kroening, D., Purandare, M., Weissenbacher, G.: Interpolant strength. In: Barthe, G., Hermenegildo, M. (eds.) VMCAI 2010. LNCS, vol. 5944, pp. 129–145. Springer, Heidelberg (2010)
8. Dutertre, B., Moura, L.D.: The Yices SMT solver. Technical report, SRI International (2006)
9. Esparza, J., Kiefer, S., Schwoon, S.: Abstraction refinement with craig interpolation and symbolic pushdown systems. In: Hermanns, H. (ed.) TACAS 2006. LNCS, vol. 3920, pp. 489–503. Springer, Heidelberg (2006)
10. Filliâtre, J.C., Marché, C.: Multi-prover verification of C programs. In: Davies, J., Schulte, W., Barnett, M. (eds.) ICFEM 2004. LNCS, vol. 3308, pp. 15–29. Springer, Heidelberg (2004)
11. Henzinger, T.A., Jhala, R., Majumdar, R., McMillan, K.L.: Abstractions from proofs. In: POPL 2004, pp. 232–244. ACM, New York (2004)
12. Jhala, R., Mcmillan, K.L.: A practical and complete approach to predicate refinement. In: Hermanns, H. (ed.) TACAS 2006. LNCS, vol. 3920, pp. 459–473. Springer, Heidelberg (2006)
13. Jhala, R., McMillan, K.L.: Array abstractions from proofs. In: Damm, W., Hermanns, H. (eds.) CAV 2007. LNCS, vol. 4590, pp. 193–206. Springer, Heidelberg (2007)
14. Jung, Y., Kong, S., Wang, B.Y., Yi, K.: Deriving invariants by algorithmic learning, decision procedures, and predicate abstraction. In: Barthe, G., Hermenegildo, M. (eds.) VMCAI 2010. LNCS, vol. 5944, pp. 180–196. Springer, Heidelberg (2010)
15. Kong, S., Jung, Y., David, C., Wang, B.Y., Yi, K.: Automatically inferring quantified loop invariants by algorithmic learning from simple templates. In: Ueda, K. (ed.) APLAS 2010. LNCS, vol. 6461, pp. 328–343. Springer, Heidelberg (2010)
16. Kroening, D., Strichman, O.: *Decision Procedures an algorithmic point of view*. EATCS. Springer, Heidelberg (2008)
17. McMillan, K.L.: An interpolating theorem prover. *Theoretical Computer Science* 345(1), 101–121 (2005)
18. McMillan, K.L.: Lazy abstraction with interpolants. In: Ball, T., Jones, R.B. (eds.) CAV 2006. LNCS, vol. 4144, pp. 123–136. Springer, Heidelberg (2006)
19. McMillan, K.L.: Quantified invariant generation using an interpolating saturation prover. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 413–427. Springer, Heidelberg (2008)
20. Srivastava, S., Gulwani, S.: Program verification using templates over predicate abstraction. In: PLDI, pp. 223–234. ACM, New York (2009)