



# Accelerating Syntax-Guided Invariant Synthesis

Grigory Fedyukovich<sup>1</sup>(✉)  and Rastislav Bodík<sup>2</sup>

<sup>1</sup> Princeton University, Princeton, USA  
grigoryf@cs.princeton.edu

<sup>2</sup> University of Washington, Seattle, USA  
bodik@cs.washington.edu

**Abstract.** We present a fast algorithm for syntax-guided synthesis of inductive invariants which combines enumerative learning with inductive-subset extraction, leverages counterexamples-to-induction and interpolation-based bounded proofs. It is a variant of a recently proposed probabilistic method, called `FREQHORN`, which is however less dependent on heuristics than its predecessor. We present an evaluation of the new algorithm on a large set of benchmarks and show that it exhibits a more predictable behavior than its predecessor, and it is competitive to the state-of-the-art invariant synthesizers based on Property Directed Reachability.

## 1 Introduction

Syntax-guided techniques [1] recently earned significant success in the field of synthesis of inductive invariants [13] for a given program and a given safety specification. Invariants are needed to represent over-approximations of the set of reachable program states, such that from their empty intersection with the set of error states one could conclude that the program is safe. While searching for invariants, it is intuitive to collect various statistics from the syntactical constructions, which appear in the program's source code, and use them as a guidance.

This work continues the track of `FREQHORN`, a completely automatic approach for (1) construction of the formal grammar based on the symbolic program encoding, and (2) probabilistic search through the candidate formulas belonging to that grammar. `FREQHORN` utilizes an SMT solver for checking inductiveness of each generated formula and iteratively constructs a suitable invariant based on the successful attempts (those formulas are called lemmas). Since based on a finite number of expressions, the formal grammar is sufficiently small, and thus the candidate formulas can be enumerated relatively quickly. We distinguish two types of candidates: (1) formulas directly extracted from the program's encoding (called *seeds*) and (2) formulas which are syntactically close to seeds (called *mutants*).

The conceptual novelty of `FREQHORN` is believed to be in the combined use of seeds and mutants, but the original paper [13] is largely silent on the

matter. Furthermore, it turns a blind eye to some algorithmic and practical details which are required for making the approach actually efficient. Among the downsides are (1) the treatment of all syntactic expressions equally and ignorance to whether the candidates have any relevance to the given safety specification; and (2) inability to predict a more-or-less appropriate order of candidates to be sampled and checked.

Luckily, elements of the Property Directed Reachability (PDR) [4, 10] can be adapted in various stages of FREQHORN’s workflow and can mitigate the downsides of the original algorithm. In particular, we propose to check candidates in batches, and we show that in practice it helps discovering larger amounts of lemmas. Additionally, we propose to keep a history of counterexamples-to-induction (CTI) which blocked FREQHORN from learning a lemma. With some periodicity, our new algorithm checks if there is a CTI which is invalidated by the currently learned lemmas, and this triggers the re-check of that failed lemma.

Last but not least, we integrate our new algorithm with the classic techniques based on Bounded Model Checking [3]. We propose to compute additional candidates by Craig interpolation [6] from proofs of bounded safety. We show that it is often sufficient to obtain some fixed amount of candidates from interpolants in the beginning of the synthesis process, and further to *bootstrap* the initial set of learned lemmas by the inductive subset extracted from the combination of the syntactic seeds and interpolants. In contrast to the entirely randomized workflow of the original version of FREQHORN, the behavior of our revised implementation at the bootstrapping is predictable. The randomized search is used by the new algorithm only for discovering mutants; and in our experiments, it was required in about one third of cases only.

To sum up, the paper contributes to the previous knowledge in the following main respects:

- A new revision and a new implementation of the FREQHORN algorithm which is split into the bootstrapping and the sampling stages. In the first stage, it deterministically exploits seeds only. In the second stage, it keeps generating and checking only mutants, and it is by design nondeterministic.
- In the bootstrapping stage, interpolation-based proofs of bounded safety that replenish the set of seeds by the candidates that likely reflect the nature of the error unreachability and consequently affect the grammar-based generation of mutants.
- In the sampling stage, the routine to extract inductive subsets which mitigates the effect of an unpredictably chosen sampling order.
- A more accurate strategy for the search space pruning and an efficient counterexample-guided approach to give some failed candidates a second chance.

The rest of the paper is structured as follows. In Sect. 2, we briefly formulate the inductive synthesis problem, and in Sect. 3 we sketch the basic FREQHORN algorithm that attempts to solve it. With the help of techniques from Sect. 4, in Sect. 5 the FREQHORN algorithm gets augmented and reformulated. In Sect. 6,

we show the experimental evidence that it indeed outperforms its predecessor and is competitive to state-of-the-art. Finally, the related work, conclusion, and acknowledgments complete the paper in Sects. 7 and 8.

## 2 Background and Notation

A first-order theory  $\mathcal{T}$  consists of a signature  $\Sigma$ , which gathers variables, function and predicate symbols, and a set  $Expr$  of  $\Sigma$ -formulas. Formula  $\varphi \in Expr$  is called  $\mathcal{T}$ -satisfiable if there exists an interpretation  $m$  of each element (i.e., a variable, a function or a predicate symbol), under which  $\varphi$  evaluates to  $\top$  (denoted  $m \models \varphi$ ); otherwise  $\varphi$  is called  $\mathcal{T}$ -unsatisfiable (denoted  $\varphi \implies \perp$ ). The Satisfiability Modulo Theory (SMT) problem [8] for a given theory  $\mathcal{T}$  and a formula  $\varphi$  aims at determining whether  $\varphi$  is  $\mathcal{T}$ -satisfiable. In this work, we formulate the tasks arising in program verification by encoding them to the SMT problems.

**Definition 1.** A transition system  $P$  is a tuple  $\langle V \cup V', Init, Tr \rangle$ , where  $V'$  is a primed copy of a set of variables  $V$ ;  $Init$  and  $Tr$  are  $\mathcal{T}$ -encodings of respectively the initial states and the transition relation.

We view *programs* as *transition systems* and throughout the paper use both terms interchangeably. *Verification task* is a pair  $\langle P, Bad \rangle$ , where  $P = \langle V \cup V', Init, Tr \rangle$  is a program, and  $Bad$  is a  $\mathcal{T}$ -encoding the *error states*. A verification task has a solution if the set of error states is unreachable. A solution to the verification task is represented by a *safe inductive invariant*, a formula that covers every initial state, is closed under the transition relation, and does not cover any of the error states.

**Definition 2.** Let  $P = \langle V \cup V', Init, Tr \rangle$ ; a formula  $Inv$  is a safe inductive invariant if the following conditions (respectively called an *initiation*, a *consecution*, and a *safety*) hold:

$$Init(V) \implies Inv(V) \tag{1}$$

$$Inv(V) \wedge Tr(V, V') \implies Inv(V') \tag{2}$$

$$Inv(V) \wedge Bad(V) \implies \perp \tag{3}$$

To simplify reading, in the rest of the paper *safe inductive invariants* are referred to as just *invariants*. We assume that an invariant  $Inv$  has the form of conjunction, i.e.,  $Inv = \ell_0 \wedge \dots \wedge \ell_n$ , and each  $\ell_i$  is called a *lemma*.

The validity of each implication (1) and (2) is equivalent to the unsatisfiability of the negation of the corresponding formula. Suppose, a formula  $Inv$  makes (1) valid, but does not make (2) valid. Thus, there exists an interpretation  $m$  satisfying  $Inv(V) \wedge Tr(V, V') \wedge \neg Inv(V')$ , to which we refer to as a *counterexample-to-induction* (CTI).

*Example 1.* The loop in program in Fig. 1a iterates  $N$  times, and in each iteration it nondeterministically picks a value  $M$ , adds it to  $x$  (conditionally) and to  $c$ ,

<pre> int x, k, c = 0; int N = NONDET(); while (c &lt; N) {   int M = NONDET();   if (k mod 2 == 0)     x = x + M;   c = c + M;   k = x + c; } assert (x ≥ N); </pre>	<pre> x = 0 c = 0 k = 0 c &lt; N k mod 2 = 0 k = x + c c ≥ N x ≥ N </pre>	<pre> α ::= 1   -1 β ::= 0   2 γ ::= x   y   k   N δ ::= α · γ + ... + α · γ   γ mod β cand ::= δ = β   δ &gt; β   δ ≥ β </pre>
(a)	(b)	(c)

$$\left\{ \begin{array}{l}
x = 0 \wedge k = 0 \wedge c = 0 \implies \mathbf{Inv}(x, k, c, N) \\
\mathbf{Inv}(x, k, c, N) \wedge c < N \wedge x' = \mathbf{ite}(k \bmod 2 = 0, x + M, x) \wedge \\
\quad c' = c + M \wedge k' = x' + c' \implies \mathbf{Inv}(x', k', c', N) \\
\mathbf{Inv}(x, k, c, N) \wedge c \geq N \wedge \neg(x \geq N) \implies \perp
\end{array} \right.$$

(d)

**Fig. 1.** Loopy program (a), its encoding (d), subexpressions extracted from the encoding (b), and grammar that generalizes the subexpressions (c).

and assigns the sum of  $x$  and  $c$  to  $k$ . We wish to prove that after the loop terminates,  $x \geq N$ . An invariant for the program is defined non-uniquely, e.g., both the conjunction  $(k \bmod 2 = 0 \wedge x = c)$  and conjunction  $(k = x + c \wedge x \geq c)$  are the solutions for this verification task.  $\square$

### 3 Syntax-Guided Invariant Synthesis

In this work, we aim at discovering invariants in an enumerative way, i.e., by guessing a candidate formula, substituting it for conditions (1), (2), and (3), and checking their validity. Here we present a moderately reformulated and simplified view of an algorithm recently proposed in [13].<sup>1</sup> The pseudocode of the algorithm, called **FREQHORN**, is shown in Algorithm 1. The key insight behind the algorithm is the automatic construction of a grammar  $G$  (line 2) based on a

<sup>1</sup> The original description [13] focuses on the probabilistic routines. In the interest of this work, we do not discuss them here but restrict our attention on describing and exemplifying the pre-processing steps.

---

**Algorithm 1.** FREQHORN: Sampling inductive invariants, cf. [13].
 

---

**Input:**  $\langle P, Bad \rangle$ : verification task, where  $P = \langle V \cup V', Init, Tr \rangle$ 
**Output:**  $Lemmas \subseteq 2^{Expr}$ 

```

1  $Seeds \leftarrow GETSUBEXPRS(Init, Tr, Bad)$ ;
2  $G \leftarrow GETGRAMMARANDDISTRIBUTIONS(Seeds)$ ;
3  $Lemmas \leftarrow \emptyset$ ;
4 while  $Bad(V) \wedge \bigwedge_{\ell \in Lemmas} \ell(V) \not\Rightarrow \perp$  do
5    $cand \leftarrow SAMPLE(G)$ ;
6    $res \leftarrow (Init(V) \Rightarrow cand(V)) \wedge (cand(V) \wedge \bigwedge_{\ell \in Lemmas} \ell(V) \wedge Tr(V, V') \Rightarrow cand(V'))$ ;
7   if  $res$  then
8      $Lemmas \leftarrow Lemmas \cup \{cand\}$ ;
9    $G \leftarrow ADJUST(G, cand, res)$ ;
```

---

fixed set  $Seeds$  of expressions obtained by traversing parse trees of  $Init$ ,  $Tr$ , and  $Bad$  (line 1).

To create  $G$  from  $Seeds$ , we drop all expressions that contain variables from both,  $V$  and  $V'$ , and deprime all variables in the remaining expressions. Then, we normalize elements of  $Seeds$  to have the form of equalities, inequalities, or disjunctions of equalities and inequalities. Finally, formulas are rewritten, such that all terms are moved to the left side, and the subtraction,  $<$ , and  $\leq$  are rewritten respectively as the addition,  $>$ , and  $\geq$ .

The algorithm uses  $G$  for generating the candidate formulas (line 5) and populates the set of lemmas until their conjunction is an invariant. The algorithm learns from each positive and negative attempt (line 9). That is,  $G$  gets adjusted, such that the candidate (and some of its close relatives) is not going to be sampled in any of the following iterations.

*Example 2.* The verification condition for the program in Fig. 1a is represented by three implications in Fig. 1d. They are syntactically split into the set  $Seeds$  of expressions over  $V$ , elements of which used to contain only primed or only unprimed variables (shown in Fig. 1b). In particular, equalities  $x' = x + M$  and  $c' = c + M$  are excluded from  $Seeds$ , and equality  $k' = x' + c'$  is rewritten to  $k = x + c$ . The grammar containing the normalized expressions from  $Seeds$  is shown in Fig. 1c. It is easy to see that all lemmas consisting in both invariants from Example 1,  $(k \bmod 2 = 0 \wedge x = c)$  and  $(k = x + c \wedge x \geq c)$ , can be generated by applying the grammar's production rules recursively.  $\square$

**Definition 3.** Each formula contained in set  $Seeds$ , which is used for constructing grammar  $G$  (in line 2), is called a seed. Formula  $cand$  produced by  $G$  is called a mutant if  $cand \notin Seeds$ .

The main downside of Algorithm 1 is that it is hard to choose a sampling order for each individual lemma at the final invariant. Suppose,  $cand = (x = c)$  is sampled and checked in the first iteration of Algorithm 1. Consequently, condition (2) is not fulfilled, and it is witnessed by the following CTI:  $[x \leftarrow 0; k \leftarrow 1; c \leftarrow 0; N \leftarrow 10; x' \leftarrow 0; k' \leftarrow 7; c \leftarrow 7; M \leftarrow 7]$ . The grammar is then adjusted, such that  $x = c$  (and some syntactically relevant, stronger or equivalent formulas, e.g.  $-x = -c$ ) do not belong to the grammar anymore.

Suppose, in the second iteration of Algorithm 1,  $cand = (k \bmod 2 = 0)$ . It passes checks (1) and (2), gets inserted to set *Lemmas*, and thus it is going to be taken into account in the following iterations (see implications in lines 4 and 6). The grammar is then adjusted again, such that  $k \bmod 2 = 0$  (and some weaker or equivalent formulas, e.g.  $k \bmod 2 \geq 0$ ) do not belong to the grammar anymore. Note that if in the third iteration  $cand = (x = c)$  was sampled again, the algorithm would terminate. However, it is impossible since the sampling grammar was adjusted after both negative and positive attempts.

The opposite sampling order (i.e.,  $cand = (k \bmod 2 = 0)$  first, and  $cand = (x = c)$  then) would lead to a faster convergence of the algorithm. Since it is hard to decide which order to choose, the production rules are equipped with probability distributions that allow both orders under certain probabilities. In this paper, we propose to use a strategy which is less dependent on an order – to check candidates in batches – and we describe it in Sect. 5 in more detail.

## 4 Old Friends Are Best

In this section, we rehash two ideas widely used in symbolic model checking that can be adapted to accelerate syntax-guided invariant synthesis.

### 4.1 Interpolation-Based Proofs of Bounded Safety

Bounded Model Checking (BMC) [3] is a formal technique, primarily used for bug finding. Given a transition system  $\langle V \cup V', Init, Tr \rangle$ , set of error states *Bad*, and a non-negative integer number  $k$ , the *BMC task* is to check if there exists a path of length  $k$  ending in an error state. The idea is to unroll *Tr*  $k$  times, conjoin it with *Init* and with the negation of *Bad*, and to check the satisfiability of the resulting formula (called a BMC formula):

$$Init(V) \wedge \underbrace{Tr(V, V') \wedge Tr(V', V'') \wedge \dots \wedge Tr(V^{(k-1)}, V^{(k)})}_k \wedge \neg Bad(V^{(k)})$$

---

**Algorithm 2.** BMCITP: Obtaining bounded proofs, cf. [24, 25].

---

**Input:**  $\langle P, Bad \rangle$ : verification task, where  $P = \langle V \cup V', Init, Tr \rangle$ ,  $k$ : bound

**Output:**  $proof \subseteq 2^{Expr}$

```

1  $unr \leftarrow \top$ ;
2 for ( $i \leftarrow k$ ;  $i > 0$ ;  $i \leftarrow i - 1$ ) do
3    $unr \leftarrow unr \wedge Tr(V^{(i-1)}, V^{(i)})$ ;
4   if  $unr \wedge Bad(V^{(k)}) \implies \perp$  then
5      $proof \leftarrow \text{GETITP}(unr, Bad(V^{(k)}))$ ;
6     return;
7  $unr \leftarrow unr \wedge Bad(V^{(k)})$ ;
8 if  $Init(V^{(0)}) \wedge unr \implies \perp$  then
9    $proof \leftarrow \text{GETITP}(Init(V^{(0)}), unr)$ ;

```

---

Here, each  $V^{(i)}$  is a fresh copy of  $V$ . Each satisfying assignment to the BMC formula represents a counterexample of length  $k$ . Otherwise, if the formula is unsatisfiable, then no counterexample of length  $k$  exists.

**Lemma 1.** *If a BMC formula for program  $P$  and some  $k$  is satisfiable then no invariant exists.*

A proof of *bounded safety* is an over-approximation  $I$  of the set of initial states, such that any path of length  $k$ , that starts in a state satisfying  $I$ , does not end in a state satisfying  $Bad$ . Extraction of proofs is typically done with the help of Craig interpolation [6].

**Definition 4.** *Given two formulas  $A$  and  $B$ , such that  $A \wedge B \implies \perp$ , an interpolant  $I$  is a formula satisfying three conditions: (1)  $A \implies I$ , (2)  $I \wedge B \implies \perp$ , and (3)  $I$  is expressed over the common alphabet to  $A$  and  $B$ .*

For an invocation of a procedure of generating an interpolant  $I$  for  $A$  and  $B$  and splitting it to a set of conjunction-free clauses (i.e.,  $I = \ell_0 \wedge \dots \wedge \ell_n$ ), we write  $\{\ell_i\} \leftarrow \text{GETITP}(A, B)$ . Algorithm 2 shows an algorithm to generate interpolation-based proofs of bounded safety for BMC formulas. It iteratively unrolls the transition relation and applies the interpolation to the entire BMC formula. In addition, in spirit of Lazy Annotation [25], while decrementing  $i$ , the algorithm applies a backward reasoning and checks if an error state is reachable by  $(k - i)$  steps from an empty state (line 4). It triggers the interpolation to be applied to smaller formulas, and in some cases fastens the proof search (line 5).

*Example 3.* Let the program in Fig. 1a is unrolled 0 times, then its BMC formula is constructed as follows:  $\underbrace{x = 0 \wedge k = 0 \wedge c = 0}_{Init} \wedge \underbrace{c \geq N \wedge \neg(x \geq N)}_{Bad}$ . It is

unsatisfiable, and since interpolants are not unique, function  $\text{GETITP}(Init, Bad)$  could return  $proof_1 = \{x \geq 0, c \leq 0\}$ ,  $proof_2 = \{x = c\}$ , or  $proof_3 = \{x \geq c\}$ .  $\square$

---

**Algorithm 3.** HOUDINI: Calculating an inductive subset, cf. [14] and keeping counterexamples-to-induction.

---

**Input:**  $P = \langle V \cup V', \text{Init}, \text{Tr} \rangle$ : program;  $Cands \subseteq 2^{Expr}$ ;  
 $CTI \subseteq 2^{V \rightarrow \mathbb{R}}$ ;  $CTI\text{map}: CTI \rightarrow 2^{Expr}$

**Output:** inductive  $Cands \subseteq 2^{Expr}$ ; updated  $CTI$  and  $CTI\text{map}$

```

1 while  $\bigwedge_{cand' \in Cands} cand'(V) \wedge Tr(V, V') \not\Rightarrow \bigwedge_{cand' \in Cands} cand'(V')$  do
2   for  $cand \in Cands$  do
3     if  $\exists \pi$ , s.t.  $\pi \models (\bigwedge_{cand' \in Cands} cand'(V) \wedge Tr(V, V') \wedge \neg cand(V'))$  then
4        $Cands \leftarrow Cands \setminus \{cand\}$ ;
5        $CTI \leftarrow CTI \cup \{\pi|_V\}$ ;
6        $CTI\text{map}(\pi|_V) \leftarrow CTI\text{map}(\pi|_V) \cup cand$ ;
```

---

## 4.2 Inductive Subset Extraction

When checking the consecution of a set of candidate formulas “one-by-one” (i.e., like in Algorithm 1), the order of checks is crucial, and the chance to miss some important lemma is high. It can be overcome by checking all candidate formulas at once, identifying which ones brake the validity of implication (2), removing them from the set, and repeating the “all-at-once” check. Algorithm 3 shows a simple implementation of this iterative algorithm, which is extensively used in PDR and also known as HOUDINI [14], Note that HOUDINI is only meaningful for the candidate formulas which are already implied by the initial states.

*Example 4.* Conjunction of formulas from set *Seeds* in Fig. 1b is unsatisfiable, and its minimal unsatisfiable core is  $c < N \wedge c \geq N$ . Thus, Algorithm 3 would immediately return the entire set *Seeds*. Let a set *Cands* be constructed from *Seeds* by removing all elements, for which condition (1) does not hold. Conjunction of the elements in *Cands* is satisfiable:  $\{x = 0, c = 0, k = 0, k \bmod 2 = 0, k = x + c\}$ . Applying Algorithm 3 to *Cands* gives the inductive subset  $\{k \bmod 2 = 0, k = x + c\}$ .  $\square$

Note that we extended Algorithm 3 with a routine to extract a counterexample-to-induction  $\pi$  for each element dropped from *Cands* (lines 3–6). We restrict each  $\pi$  to only assignments to variables from  $V$  (denoted  $\pi|_V$ ) and group all non-inductive formulas from *Cands* by the particular  $\pi$  that killed them. This routine is important for optimizing the syntax-guided invariant synthesis algorithm, and it is discussed in more detail in Sect. 5.

## 5 Reconsidering Syntax-Guided Invariant Synthesis

The lesson we learned when running the FREQHORN algorithm is that the program encoding gives many hints on how the shape of lemmas should look like.



However, the encoding itself can barely give any information about the sampling order. Our main idea to revise the FREQHORN algorithm is to treat seeds and mutants separately. Indeed, as we have seen in Example 2, both seeds and mutants are needed for constructing an invariant, but seeds do *not* actually need to be re-sampled – these candidates are ready to be checked prior to any sampling.

## 5.1 Overview

We present a new revision of the FREQHORN algorithm which is split into two main stages, the *bootstrapping* and the *sampling*. In the first stage, it exploits only seeds. The idea is to terminate this stage as quickly as possible and to populate the set of lemmas with (preferably, the maximal) inductive subset of seeds. If this subset is not enough for an invariant, the algorithm should proceed to the next stage, in which it should keep generating and checking only mutants.

The pseudocode of the new FREQHORN’s revision is shown in Algorithm 4. In the bootstrapping, the algorithm relies on Algorithm 2 to replenish the set of seeds by semantically-meaningful candidates, and in the sampling stage, it relies on Algorithm 3 to mitigate the effect of an unpredictably chosen sampling order. Another algorithmic advantage against Algorithm 1 (to be explained in Sect. 5.2) lies in a more accurate strategy for the search space pruning and the efficient counterexample-guided method to give some failed candidates a *second chance*.

The algorithm takes as input a verification task and values of important configuration parameters  $N$ ,  $M$ , and  $K$  (to be explained further). Like Algorithm 1, it starts with obtaining a set of expressions *Seeds* from *Init*, *Tr*, and *Bad* (line 1). Then, *Seeds* gets merged with sets of formulas obtained by Craig interpolation from proofs of bounded safety for a range of bounds  $0, \dots, N$ . Note that if there is a counterexample of length  $k < N$  discoverable by the BMC engine then an invariant does not exist (recall Lemma 1), and Algorithm 4 terminates (line 5).

The bootstrapping ends when the merged set *Seeds* is taken as input by Algorithm 3, and it extracts an inductive subset (line 19). However, prior to it, the algorithm checks the initiation condition for all elements of the merged set, and the set is filtered accordingly (lines 15–18).

---

**Algorithm 4.** FREQHORN-2: Sampling inductive invariants with HOUDINI, BMCITP, and the second-chance candidates.

---

**Input:**  $\langle P, Bad \rangle$ : verification task, where  $P = \langle V \cup V', Init, Tr \rangle$ ;  $N, M, K$ : knobs  
**Output:**  $Lemmas \subseteq 2^{Expr}$

```

1   $Seeds \leftarrow GETSUBEXPRS(Init, Tr, Bad)$ ;
2  for  $(k \leftarrow 0; k < N; k \leftarrow k + 1)$  do
3     $proof \leftarrow BMCITP(\langle P, Bad \rangle, k)$ ;
4    if  $proof = \emptyset$  then
5      return;
6    else
7       $Seeds \leftarrow Seeds \cup proof$ ;
8   $G \leftarrow GETGRAMMARANDDISTRIBUTIONS(Seeds)$ ;
9   $CTI, CTImap \leftarrow \emptyset$ ;
10  $\#learned \leftarrow 0$ ;
11  $Cands \leftarrow Seeds$ ;
12 while  $Bad(V) \wedge \bigwedge_{\ell \in Lemmas} \ell(V) \not\Rightarrow \perp$  do
13   while  $|Cands| < M$  do
14      $Cands \leftarrow Cands \cup \{SAMPLE(G)\}$ ;
15   for  $cand \in Cands$  do
16     if  $Init(V) \not\Rightarrow cand(V)$  then
17        $G \leftarrow ADJUST(G, cand, false)$ ;
18        $Cands \leftarrow Cands \setminus \{cand\}$ ;
19    $\langle Lemmas', CTI, CTImap \rangle \leftarrow HOUDINI(P, Cands \cup Lemmas, CTI, CTImap)$ ;
20   for  $\ell' \in Lemmas' \setminus Lemmas$  do
21      $G \leftarrow ADJUST(G, \ell', true)$ ;
22    $NewLemmas = \{\ell' \mid \ell' \in Lemmas', s.t. \bigwedge_{\ell \in Lemmas} \ell \not\Rightarrow \ell'\}$ ;
23    $\#learned \leftarrow \#learned + |NewLemmas|$ ;
24    $Lemmas = Lemmas \cup NewLemmas$ ;
25    $Cands \leftarrow \emptyset$ ;
26   if  $\#learned > K$  then
27      $\#learned \leftarrow 0$ ;
28     for  $m \in CTI$  do
29       if  $m \not\models \bigwedge_{\ell \in Lemmas} \ell(V)$  then
30          $CTI \leftarrow CTI \setminus m$ ;
31          $Cands \leftarrow Cands \cup CTImap(m)$ ;

```

---

*Example 5.* Let set  $Seeds$  be as in Fig. 1b, and set  $Cands$  be constructed from  $Seeds$  by removing all elements, for which condition (1) does not hold. Assume that a proof of bounded safety for  $k = 0$  is  $\{x = c\}$  (as one of the options in Example 3). Applying Algorithm 3 to  $Cands \cup \{x = c\}$ , we get the inductive subset  $\{k \bmod 2 = 0, k = x + c, x = c\}$ . Since the conjunction of these lemmas is an invariant, the algorithm terminates just after the bootstrapping.  $\square$

Checking the candidate formulas in batches is an important improvement over Algorithm 1. This way, the algorithm becomes less dependent of the heuristics for prioritizing the search-space traversal. The size of the batch  $M$  is configurable, and if the size of set  $Cands$  is less than  $M$ , then the set gets additional mutants (lines 13–14). Mutants are sampled from the grammar, which is powered by both, the program’s encoding (similar to Algorithm 1) and the proofs of bounded safety (new in Algorithm 4). This enlarges the search space for the further mutants.

If the initial batch of candidates still misses some lemmas necessary for an invariant, then Algorithm 4 proceeds to a new iteration. In particular, the extracted inductive subset gets merged with the set of lemmas (line 24), and the assembly of a new batch of candidates starts from scratch (line 25).

*Example 6.* Assume that a proof of bounded safety is  $proof_1 = \{x \geq 0, c \leq 0\}$  (as in Example 3). However, the initiation condition is fulfilled for none of the elements of  $proof_1$ , so none of them contains in the set of formulas  $Cands$  taken as input by Algorithm 3. Thus,  $proof_1$  does not bring any additional value to the set of seeds, and (contrary to the case in Example 5) the algorithm does not terminate after the bootstrapping. Instead, it proceeds to sampling fresh mutants.  $\square$

**Theorem 1.** *If Algorithm 4 terminates, then either an actual bug is found (line 5), or an invariant is synthesized (after the while-loop).*

## 5.2 Learning Strategy

A substantial distinction between the FREQHORN’s revisions is how they react to the positive and negative attempts. In Algorithm 1, the search space gets adjusted after each individual check (recall the example after Definition 3). The grammar adjustments are performed by changing the probabilities assigned to the production rules. In addition to zeroing the probability of sampling a candidate  $cand$  itself, after each positive check, Algorithm 1 zeroes the probabilities of sampling some formulas which are weaker than  $cand$ , and after each negative check – the probabilities of sampling some formulas which are stronger than  $cand$  (see [13] for more details).

In contrast, Algorithm 4 reacts just to the failed candidates after the initiation check (line 17) and to the successful candidates after the consecution check (line 21). Otherwise, if the consecution check failed for a candidate  $cand$  (inside Algorithm 3), Algorithm 4 does not disqualify  $cand$  from being checked again in the future, and this is done by keeping  $cand$  locally and periodically seeking an opportunity to give  $cand$  a *second chance*.

To efficiently exploit the second-chance candidates, we rely on the extension of Algorithm 3 by the routine to extract counterexamples-to-induction. That is, for each failed  $cand$  there exists  $m \in CTI$  that killed it. To maintain this information, every application of Algorithm 3 updates the map  $CTImap$  from  $CTI$  to failed candidates. In Algorithm 4, it remains to periodically check whether

some  $m$  is eliminated (line 29), and it would increase chances of all candidates killed by  $m$  (line 31) to succeed the consecution check in the next iteration. On the other hand, if some  $m$  still models the conjunction of learned lemmas then it is guaranteed that candidates in  $CTImap(m)$  will fail the consecution check again.

Finally, to ensure that the CTI-check happens not too often, we run it only when at least  $K$  new lemmas are learned. To make this happen, Algorithm 4 performs a redundancy check (line 22) for all lemmas  $\ell$  that have passed the initiation and the consecution checks:  $\ell$  gets learned only when the conjunction of all lemmas learned so far does not imply  $\ell$ . Obviously, when no new lemmas (after the redundancy check) are added, it does not make sense to run the CTI-check since all CTIs are still valid.

### 5.3 Optimizations

The following tricks are omitted from the algorithm’s pseudocode to simplify reading, but they are important for the algorithm’s efficiency.

- As a consequence of calculating frequencies, in the original FREQHORN algorithm, seeds were given priorities, but mutants were considered with a relatively small probability. In contrast, the new FREQHORN’s revision forces seeds to be checked in the bootstrapping. So while doing sampling, it gives priorities to mutants, and for that it ignores frequencies.
- The initiation checks (lines 15–18) for proofs of bounded safety are omitted since by definition of interpolant (Definition 4) they are already fulfilled. The initiation checks for the second-chance candidates are omitted as well.
- In case a candidate fails the consecution check, and it is queued for a second chance, it is still possible that Algorithm 4 samples it again in the next iterations. Re-sampling is avoided by additional adjustments to the probabilities of the sampling grammar in line 6 of Algorithm 3.
- Algorithm 3 could be optimized if solved *with assumptions*. However, in our experience, it may lead to dropping more candidates than needed. Ideas for getting a maximal inductive subset from [23] could be applied here as well.
- For getting proofs of bounded safety for various bounds, an incremental SMT solver could be used. That is, it could reuse parts of a BMC formula for bound  $k$  to encode the BMC formula for bound  $k + 1$ . Potentially, other tricks (e.g., [5, 30]) could also be applied here. Finally, interpolation could be replaced by the weakest precondition computation.

## 6 Evaluation

We implemented FREQHORN-2 on top of our prior implementation FREQHORN<sup>2</sup>. The tool takes as input a verification task in a form of linear constrained Horn

---

<sup>2</sup> The source code and benchmarks are available at <https://github.com/grigoryfedukovich/aeval/tree/rnd>.

clauses (CHC), automatically performs its unrolling, searches for counterexamples, generates proofs of bounded safety, and performs the HOUDINI-style extraction of inductive subsets. All the symbolic reasoning is performed by the Z3 SMT solver [7].

We evaluated FREQHORN-2 on various safe and buggy programs taken from SVCOMP<sup>3</sup> and literature (e.g., [9, 15]). Since most of benchmarks, proposed by [9], appeared to be solvable during the bootstrapping of FREQHORN-2 (more details in Sect. 6.1) within (fractions of) seconds, we crafted additional harder benchmarks by ourselves.

All the programs were encoded using the theories of linear (LIA) and non-linear integer arithmetic (NIA). We did run FREQHORN-2 on unsafe instances for the testing purposes only. It was able to detect a counterexample, but since no invariant exists in these cases, we do not discuss this experience here.

## 6.1 The Bootstrapping Experiment

In total, we considered **171** safe programs. For **103** of them, the seeds, generated by breaking the symbolic encoding to pieces, did already contain all lemmas needed for invariants. However, when we checked the seeds one-by-one, we revealed invariants for only 63, but using the inductive subset extraction helped revealing all 103. Each set of seeds contained in average 9 formulas.

For our BMC implementation, we considered bounds 1, 2, and 3. Generated interpolants already contained all lemmas for invariants for **70** programs.<sup>4</sup> Each set of bounded proofs contained in average 2 formulas. In all these cases, the output of Algorithm 2 was taken as input by Algorithm 3, and the final safety check was performed afterwards. Our most promising results were achieved while running Algorithm 3 for the merged sets of seeds and proofs of bounded safety (i.e., both sets as in the two prior runs together). The merged sets already contained all lemmas for invariants for **114** programs.

This experiment lets us to conclude that *the bootstrapping is exceptionally important* for accelerating syntax-guided invariant synthesis. In contrast to FREQHORN’s fully randomized workflow, FREQHORN-2’s behavior at the bootstrapping is predictable. FREQHORN-2 uses the randomized search only to discover mutants, and in our experiments, it was required only in 57 out of 171 cases.

## 6.2 Overall Statistics

Since technically FREQHORN-2 is a CHC-solver, we compared it against other CHC-solvers, namely  $\mu Z$  [17, 22] and SPACER3 [22]. All the tools were provided with the same CHC-encodings of verification problems (and thus, the results do not directly depend on a process of encoding a C-program to a CHC-file). Both  $\mu Z$  and SPACER3 are PDR-based, and despite the latter is faster than the

<sup>3</sup> Software Verification Competition, <http://sv-comp.sosy-lab.org/>, loop-\* categories.

<sup>4</sup> Currently interpolation in FREQHORN-2 is limited to LIA, so we had to skip interpolation for 17 benchmarks over NIA.

Table 1. Exact timings.

Benchmark	FREQHORN.2	FREQHORN	$\mu Z$	SPACER	Benchmark	FREQHORN.2	FREQHORN	$\mu Z$	SPACER
abdu_01	ε	1.27	ε	ε	nonlin_minus_2	ε	2.60	∞	∞
abdu_04	ε	8.64	ε	ε	nonlin_mod_1	5.78	16.13	∞	∞
bouncy_three	6.87	6.76	ε	ε	nonlin_mod_2	∞	18.46	∞	∞
bouncy_two	2.74	2.41	∞	∞	nonlin_mod_mult	ε	ε	∞	∞
cegar1	ε	2.18	ε	ε	nonlin_mult_1	ε	ε	∞	∞
cggmp_iter_1	3.09	8.73	ε	ε	nonlin_mult_2	ε	2.28	∞	∞
cggmp_iter_2	10.50	24.92	ε	ε	nonlin_mult_3	ε	2.86	∞	∞
cggmp_iter_3	18.80	55.86	ε	ε	nonlin_mult_4	ε	2.48	∞	∞
const_div_1	ε	2.61	ε	ε	nonlin_mult_5	∞	3.22	∞	∞
const_div_2	ε	1.40	ε	ε	nonlin_mult_6	ε	3.27	∞	∞
const_div_3	∞	∞	21.31	∞	nonlin_power	ε	∞	ε	ε
const_mod_3	ε	∞	ε	ε	nonlin_square	ε	ε	∞	∞
count_by_2_modif	3.28	10.98	∞	∞	nonterm_01	ε	ε	∞	∞
count_by_2	2.80	1.62	∞	∞	phases1	ε	10.81	∞	∞
countud	3.60	15.96	∞	ε	s_disj_ite_01	8.77	5.43	ε	ε
css2003	ε	9.59	ε	ε	s_disj_ite_02	3.98	5.98	ε	ε
dillig02	10.91	35.48	ε	ε	s_disj_ite_03	13.15	2.71	ε	ε
dillig03	ε	ε	∞	∞	s_disj_ite_04	7.01	4.13	ε	ε
dillig05	4.74	2.04	∞	∞	s_disj_ite_05	4.64	∞	18.66	40.38
dillig07	1.93	2.58	ε	ε	s_disj_ite_06	6.43	∞	ε	ε
dillig08	ε	1.26	∞	ε	s_mutants_01	ε	ε	∞	∞
dillig10	2.28	2.72	ε	ε	s_mutants_02	ε	2.40	∞	∞
dillig13.1	∞	∞	ε	ε	s_mutants_03	ε	ε	∞	∞
dillig13	∞	∞	ε	ε	s_mutants_05	7.86	1.90	∞	∞
dillig14	1.90	6.98	ε	ε	s_mutants_06	ε	ε	∞	∞
dillig15	3.28	3.61	∞	ε	s_mutants_07	1.63	1.75	ε	ε
dillig16	20.14	15.35	∞	ε	s_mutants_09	ε	3.41	ε	ε
dillig20.1	ε	1.65	ε	ε	s_mutants_11	ε	1.86	4.72	ε
dillig20.2	ε	4.97	ε	ε	s_mutants_12	ε	4.09	ε	ε
dillig20.3	34.51	4.53	ε	ε	s_mutants_13	ε	10.42	ε	ε
dillig21	ε	4.06	7.5	3.98	s_mutants_14	ε	11.47	ε	ε
dillig22.2	ε	8.59	ε	ε	s_mutants_15	ε	14.50	ε	ε
dillig22.3	29.77	15.43	ε	ε	s_mutants_16	14.38	∞	∞	ε
dillig22.4	9.99	15.20	ε	ε	s_mutants_17	14.56	5.83	∞	∞
dillig22.5	14.12	12.02	∞	ε	s_mutants_18	39.85	∞	ε	ε
dillig22.6	25.88	13.75	∞	ε	s_mutants_19	ε	1.31	1.44	ε
dillig22	ε	7.25	ε	ε	s_mutants_20	41.15	36.47	∞	∞
dillig37	1.48	1.67	ε	ε	s_mutants_21	ε	24.68	∞	∞
dillig41	ε	1.41	∞	∞	s_mutants_22	14.82	16.52	∞	∞
dillig42.1	23.88	∞	ε	ε	s_mutants_23	1.61	1.11	∞	∞
dillig42	52.13	∞	ε	ε	s_mutants_24	1.58	2.81	∞	∞
dillig44.1	ε	∞	ε	ε	s_seeds_04	ε	ε	20.49	ε
dillig44	ε	∞	ε	ε	s_seeds_05	ε	∞	∞	ε
dillig46	6.20	4.86	ε	ε	s_seeds_06	ε	1.78	ε	∞
ez7	2.28	8.93	ε	ε	s_seeds_10	ε	ε	∞	∞
exact_iters_1	26.55	35.58	ε	ε	s_triv_01	ε	1.33	ε	ε
exact_iters_2	∞	22.76	∞	∞	s_triv_07	ε	1.15	ε	ε
exact_iters_3	23.96	26.60	∞	∞	s_triv_08	ε	∞	ε	ε
exact_iters_4	ε	29.19	ε	ε	s_triv_09	ε	∞	ε	ε
fig3	ε	21.82	ε	ε	s_triv_11	ε	∞	ε	ε
formula22	ε	11.34	ε	ε	s_triv_12	ε	∞	ε	ε
formula25	ε	1.12	ε	ε	s_triv_14	ε	∞	ε	ε
formula27	ε	20.31	ε	ε	s_triv_16	ε	∞	ε	ε
gcd_2	2.15	1.84	ε	ε	s_triv_17	ε	∞	ε	ε
gcd_3	2.40	3.58	ε	ε	sn_1024	6.14	14.64	24.83	7.12
gj2007	∞	∞	∞	ε	sn_2048	25.91	26.11	∞	30.87
half_true_modif	ε	58.15	ε	ε	sn_4096	7.57	∞	∞	∞
half_true_orig	35.58	45.00	ε	ε	sn_8192	11.82	41.01	∞	∞
hkk2008	6.60	28.46	ε	ε	three_dots_moving_1	ε	58.48	10.28	ε
menlo_park_term_simpl_1	ε	1.39	∞	15.84	three_dots_moving_2	ε	∞	ε	ε
menlo_park_term_simpl_2	12.29	58.36	∞	∞	three_dots_moving_3	ε	∞	ε	ε
n_c11	ε	∞	ε	ε	trex3	ε	8.23	ε	ε
nonlin_div	26.00	∞	∞	∞	yz_plus_minus	4.35	∞	∞	∞
nonlin_factorial	ε	∞	∞	∞					
nonlin_minus_1	ε	1.78	∞	∞					




Fig. 2. FREQHORN-2 vs respectively FREQHORN, SPACER3, and  $\mu Z$ .

former and can solve more benchmarks, there are 26 instances, for which the former outperforms the latter.

Table 1 shows the precise running times of `FREQHORN-2`, `FREQHORN`,  `$\mu$ Z`, and `SPACER3`. To simplify reading, we removed non-representative “*noise*”-runs which took less than 1 s or exceeded a timeout of 60 s by all tools. In the table,  $\epsilon$  denotes an insignificant amount of time ( $\leq 1$  s), and  $\infty$  denotes the timeout. The numbers of `FREQHORN` and `FREQHORN-2` are the means of three individual runs. In total, the table contains **128** instances. Additionally, Fig. 2 shows three scatter plots comparing running times of `FREQHORN-2` vs `FREQHORN`,  `$\mu$ Z`, and `SPACER3` respectively. Each point in a plot represents a pair of the `FREQHORN-2` run (x-axis) and the competing tool run (y-axis).

`FREQHORN-2` outperformed its predecessor in **90** out of 128 cases. We witnessed the speedup up to **233X**, and in average `FREQHORN-2` was *four times* faster than `FREQHORN`. In **40** cases `FREQHORN-2` outperformed `SPACER3`, and in 38 cases `SPACER3` outperformed `FREQHORN-2`. In **51** cases `FREQHORN-2` outperformed  `$\mu$ Z`, and in 34 cases  `$\mu$ Z` outperformed `FREQHORN-2`. Unfortunately, `FREQHORN-2` still has some performance anomalies, which we believe are connected to the often blind grammar-construction mechanism, inability to generate large disjunctions, and possible inefficiencies of the black-box interpolation engine.

## 7 Related Work

In this work we exploit a range of techniques originated from symbolic model checking, and in particular from `IC3/PDR` [4, 10], e.g., the idea of keeping CTIs and analyzing them to push previously considered lemmas [29]. Various strategies could be applied for making the lemma pushing more or less eager, i.e., as soon as a newly-added lemma invalidates some CTI. In some `IC3` implementations (e.g., [16]), eager pushing does not pay off, but avoiding to push certain lemmas during the regular pushing stage of `IC3` results in an improvement. Since we do not have many lemmas, eager pushing also pays off.

The idea of applying `HOUDINI` to extract invariants from proofs of bounded safety was fundamental for the first version of `SPACER` [23]. They, however, keep obtaining proofs along the entire verification process. In contrast, we use proofs mainly for the bootstrapping, while the remaining progress of the algorithm is entirely dictated by the success of sampling.

Most of the successful verification tools today use various combinations of different techniques. In particular, approaches [2, 28] use invariants from abstract interpretation to force convergence of k-induction. Recently, k-induction was benefitted from lemmas obtained from `PDR` [21]. A promising idea to exploit the data from traces [12, 15] while creating and manipulating the candidates for invariants could also be used in our syntax-guided approach: at least we could add more constants to the grammar. However we are currently unaware of a strategy to find meaningful constants and to avoid over-population of the grammar by too many constants. Our preliminary experiments resulted so far in the performance decrease.

Techniques for automatic construction of grammars were applied outside of formal verification, but in the domains of security analysis and dynamic test generation [18, 19]. Indeed, mutations of the input data for some program can in fact be used as new input data and therefore can increase the testing coverage.

Finally, syntax-guided techniques [1] keep being used in program synthesis more frequently than in the inductive invariant synthesis. For instance, in applications [11, 20, 26, 27] a formal grammar is additionally provided, and it is considered a part of specification. In contrast, in our application, the verification condition contains the encoding of the entire program and the safety specification, which together are enough for construction of formal grammars completely automatically. This is in fact the main driving idea behind `FREQHORN`, and it leaves us a spacious room for its further adaptations, e.g., in proving and disproving program termination, automated repair of software regressions, and security analysis.

## 8 Conclusion

We have presented the new revision of the `FREQHORN` algorithm to synthesize safe inductive invariants based on syntactic features of the source code and the proofs of bounded safety. The new algorithm contains the deterministic bootstrapping stage and the nondeterministic sampling stage, which make it more predictable than its predecessor, allows converging more frequently and in average four times faster. Similarly to most of the state-of-the-art verification techniques, our approach enjoys a tight integration with well renowned formal methods and should be treated as an example of successful interchange of ideas across application domains.

**Acknowledgments.** It is hard to underestimate the value of discussions with Alexander Ivrii, Arie Gurfinkel, Michael W. Whalen, and other attendees of the International Conference on Formal Methods in Computer-Aided Design (FMCAD 2017) which gave rise to many interesting ideas and inspired this work.

## References

1. Alur, R., Bodík, R., Juniwal, G., Martin, M.M.K., Raghothaman, M., Seshia, S.A., Singh, R., Solar-Lezama, A., Torlak, E., Udupa, A.: Syntax-guided synthesis. In: FMCAD, pp. 1–17. IEEE (2013)
2. Beyer, D., Dangl, M., Wendler, P.: Boosting  $k$ -induction with continuously-refined invariants. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9206, pp. 622–640. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-21690-4\\_42](https://doi.org/10.1007/978-3-319-21690-4_42)
3. Biere, A., Cimatti, A., Clarke, E., Zhu, Y.: Symbolic model checking without BDDs. In: Cleaveland, W.R. (ed.) TACAS 1999. LNCS, vol. 1579, pp. 193–207. Springer, Heidelberg (1999). [https://doi.org/10.1007/3-540-49059-0\\_14](https://doi.org/10.1007/3-540-49059-0_14)
4. Bradley, A.R.: SAT-based model checking without unrolling. In: Jhala, R., Schmidt, D. (eds.) VMCAI 2011. LNCS, vol. 6538, pp. 70–87. Springer, Heidelberg (2011). [https://doi.org/10.1007/978-3-642-18275-4\\_7](https://doi.org/10.1007/978-3-642-18275-4_7)



5. Cabodi, G., Camurati, P., Palena, M., Pasini, P., Vendraminetto, D.: Interpolation-based learning as a mean to speed-up bounded model checking (short paper). In: Cimatti, A., Sirjani, M. (eds.) SEFM 2017. LNCS, vol. 10469, pp. 382–387. Springer, Cham (2017). [https://doi.org/10.1007/978-3-319-66197-1\\_25](https://doi.org/10.1007/978-3-319-66197-1_25)
6. Craig, W.: Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory. *J. Symbolic Logic*, 269–285 (1957)
7. de Moura, L., Björner, N.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-78800-3\\_24](https://doi.org/10.1007/978-3-540-78800-3_24)
8. Detlefs, D., Nelson, G., Saxe, J.B.: Simplify: a theorem prover for program checking. *J. ACM* **52**(3), 365–473 (2005)
9. Dillig, I., Dillig, T., Li, B., McMillan, K.L.: Inductive invariant generation via abduction inference. In: OOPSLA, pp. 443–456. ACM (2013)
10. Eén, N., Mishchenko, A., Brayton, R.K.: Efficient implementation of property directed reachability. In: FMCAD, pp. 125–134. IEEE (2011)
11. Fedyukovich, G., Ahmad, M.B.S., Bodík, R.: Gradual synthesis for static parallelization of single-pass array-processing programs. In: PLDI, pp. 572–585. ACM (2017)
12. Fedyukovich, G., D’Iddio, A.C., Hyvärinen, A.E.J., Sharygina, N.: Symbolic detection of assertion dependencies for bounded model checking. In: Egved, A., Schaefer, I. (eds.) FASE 2015. LNCS, vol. 9033, pp. 186–201. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-46675-9\\_13](https://doi.org/10.1007/978-3-662-46675-9_13)
13. Fedyukovich, G., Kaufman, S., Bodík, R.: Sampling invariants from frequency distributions. In: FMCAD, pp. 100–107. IEEE (2017)
14. Flanagan, C., Leino, K.R.M.: Houdini, an annotation assistant for ESC/Java. In: Oliveira, J.N., Zave, P. (eds.) FME 2001. LNCS, vol. 2021, pp. 500–517. Springer, Heidelberg (2001). [https://doi.org/10.1007/3-540-45251-6\\_29](https://doi.org/10.1007/3-540-45251-6_29)
15. Garg, P., Neider, D., Madhusudan, P., Roth, D.: Learning invariants using decision trees and implication counterexamples. In: POPL, pp. 499–512. ACM (2016)
16. Gurfinkel, A., Ivrii, A.: Pushing to the top. In: FMCAD, pp. 65–72. IEEE (2015)
17. Hoder, K., Björner, N.: Generalized property directed reachability. In: Cimatti, A., Sebastiani, R. (eds.) SAT 2012. LNCS, vol. 7317, pp. 157–171. Springer, Heidelberg (2012). [https://doi.org/10.1007/978-3-642-31612-8\\_13](https://doi.org/10.1007/978-3-642-31612-8_13)
18. Höschele, M., Zeller, A.: Mining input grammars from dynamic taints. In: ASE, pp. 720–725. ACM (2016)
19. Höschele, M., Zeller, A.: Mining input grammars with AUTOGRAM. In: ICSE - Companion Volume, pp. 31–34. IEEE Computer Society (2017)
20. Inala, J.P., Polikarpova, N., Qiu, X., Lerner, B.S., Solar-Lezama, A.: Synthesis of recursive ADT transformations from reusable templates. In: Legay, A., Margaria, T. (eds.) TACAS 2017. LNCS, vol. 10205, pp. 247–263. Springer, Heidelberg (2017). [https://doi.org/10.1007/978-3-662-54577-5\\_14](https://doi.org/10.1007/978-3-662-54577-5_14)
21. Jovanovic, D., Dutertre, B.: Property-directed k-induction. In: FMCAD, pp. 85–92. IEEE (2016)
22. Komuravelli, A., Gurfinkel, A., Chaki, S.: SMT-based model checking for recursive programs. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 17–34. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-08867-9\\_2](https://doi.org/10.1007/978-3-319-08867-9_2). <https://bitbucket.org/spacer/code/branch/spacer3>
23. Komuravelli, A., Gurfinkel, A., Chaki, S., Clarke, E.M.: Automatic abstraction in SMT-based unbounded software model checking. In: Sharygina, N., Veith, H. (eds.) CAV 2013. LNCS, vol. 8044, pp. 846–862. Springer, Heidelberg (2013). [https://doi.org/10.1007/978-3-642-39799-8\\_59](https://doi.org/10.1007/978-3-642-39799-8_59)

24. McMillan, K.L.: Interpolation and SAT-based model checking. In: Hunt, W.A., Somenzi, F. (eds.) CAV 2003. LNCS, vol. 2725, pp. 1–13. Springer, Heidelberg (2003). [https://doi.org/10.1007/978-3-540-45069-6\\_1](https://doi.org/10.1007/978-3-540-45069-6_1)
25. McMillan, K.L.: Lazy annotation revisited. In: Biere, A., Bloem, R. (eds.) CAV 2014. LNCS, vol. 8559, pp. 243–259. Springer, Cham (2014). [https://doi.org/10.1007/978-3-319-08867-9\\_16](https://doi.org/10.1007/978-3-319-08867-9_16)
26. Phothilimthana, P.M., Jelvis, T., Shah, R., Totla, N., Chasins, S., Bodík, R.: Chlorophyll: synthesis-aided compiler for low-power spatial architectures. In: PLDI, pp. 396–407. ACM (2014)
27. Pu, Y., Bodík, R., Srivastava, S.: Synthesis of first-order dynamic programming algorithms. In: OOPSLA, pp. 83–98. ACM (2011)
28. Roux, P., Delmas, R., Garoche, P.: SMT-AI: an abstract interpreter as oracle for k-induction. *Electr. Notes Theor. Comput. Sci.* **267**(2), 55–68 (2010)
29. Suda, M.: Triggered clause pushing for IC3. CoRR, abs/1307.4966 (2013)
30. Vizel, Y., Gurfinkel, A., Malik, S.: Fast interpolating BMC. In: Kroening, D., Păsăreanu, C.S. (eds.) CAV 2015. LNCS, vol. 9206, pp. 641–657. Springer, Cham (2015). [https://doi.org/10.1007/978-3-319-21690-4\\_43](https://doi.org/10.1007/978-3-319-21690-4_43)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

