



Ilinva: Using Abduction to Generate Loop Invariants

Mnacho Echenim, Nicolas Peltier, Yanis Sellami

► To cite this version:

Mnacho Echenim, Nicolas Peltier, Yanis Sellami. Ilinva: Using Abduction to Generate Loop Invariants. Frontiers of Combining Systems - 12th International Symposium, 2019, London, United Kingdom. pp.77-93, 10.1007/978-3-030-29007-8_5 . hal-02323446

HAL Id: hal-02323446

<https://hal.science/hal-02323446>

Submitted on 5 Nov 2019

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Ilinva: Using Abduction to Generate Loop Invariants

Mnacho Echenim, Nicolas Peltier, and Yanis Sellami

Univ. Grenoble Alpes, CNRS, LIG, F-38000 Grenoble France
`[Mnacho.Echenim|Nicolas.Peltier|Yanis.Sellami]@univ-grenoble-alpes.fr`

Abstract. We describe a system to prove properties of programs. The key feature of this approach is a method to automatically synthesize inductive invariants of the loops contained in the program. The method is generic, *i.e.*, it applies to a large set of programming languages and application domains; and lazy, in the sense that it only generates invariants that allow one to derive the required properties. It relies on an existing system called GPID for abductive reasoning modulo theories [14], and on the platform for program verification Why3 [16]. Experiments show evidence of the practical relevance of our approach.

1 Introduction

Hoare logic – together with strongest post-conditions or weakest pre-conditions calculi – allow one to verify properties of programs defined by bounded sequences of instructions [20]. Given a pre-condition ϕ satisfied by the inputs of program P , algorithms exist to compute the strongest formula ψ such that $\phi \{P\} \psi$ holds, meaning that if ϕ holds initially then ψ is satisfied after P is executed, and any formula ψ' that holds after P is executed is such that $\psi \models \psi'$. To check that the final state satisfies some formula ψ' , we thus only have to check that ψ' is a logical consequence of ψ . However, in order to handle programs containing loops, it is necessary to associate each loop occurring within the program with an *inductive invariant*. An inductive invariant for a given loop L is a formula that holds every time the program enters L (*i.e.*, it must be a logical consequence of the preconditions of L), and is preserved by the sequence of instructions in L . Testing whether a formula is an inductive invariant is a straightforward task, and the difficulty resides in generating candidate invariants. These can be supplied by the programmer, but this is a rather tedious and time-consuming task; for usability and scalability, it is preferable to generate those formulas automatically when possible. In this paper, we describe a system to generate such invariants in a completely automated way, via abductive reasoning modulo theories, based on the methodology developed in [13]. Roughly speaking, the algorithm works as follows. Given a program P decorated with a set of assertions that are to be established, all loops are first assigned the same candidate invariant \top . These invariants are obviously sound: they hold before the loops and are preserved by the sequence of instructions in the loop; however they are usually not strong enough

to prove the assertions decorating the program. They are therefore strengthened by adding hypotheses that are sufficient to ensure that the assertions hold; these hypotheses are generated by a tool that performs *abductive inferences*, and the strengthened formulas are *candidate invariants*. Additional strengthening steps are taken to guarantee that these candidates are actual invariants, *i.e.*, that they are preserved by the sequence of instructions in the loop. These steps are iterated until a set of candidate invariants that are indeed inductive is obtained.

We rely on two existing systems to accomplish this task. The first one is Why3 (see, *e.g.*, <http://why3.lri.fr/> or [16]), a well-known and widely-used platform for deductive program verification that is used to compute verification conditions and verify assertions. The second system, GPiD, is designed to generate implicants¹ of quantifier-free formulas modulo theories [14]. This system is used as an abductive reasoning procedure, thanks to the following property: if $\phi \not\models \psi$, finding a hypothesis ϕ' such that $\phi \wedge \phi' \models \psi$ is equivalent to finding ϕ' such that $\phi' \models \phi \Rightarrow \psi$. GPiD is generic, since it only relies on the existence of a decision procedure for the considered theory (counter-examples are exploited when available to speed-up the generation of the implicants when available). Both systems are connected in the ILINVA framework.

Related Work. A large number of different techniques have been proposed to generate loop invariants automatically, especially on numeric domains [9,10], but also in more expressive logics, for programs containing arrays or expressible using combination of theories [26,8,23,18,22,24]. We only briefly review the main ideas of the most popular and successful approaches. Methods based on abstract interpretations (see, *e.g.*, [11,25]) work by executing the program in a symbolic way, on some abstract domain, and try to compute over-estimations of the possible states of the memory after an unbounded number of iterations of the loop. Counter-examples generated from runs can be exploited to refine the considered abstraction [17,19]. The idea is that upon detection of a run for which the assertion is violated, if the run does not correspond to a concrete execution path, then the considered abstraction may be refined to dismiss it.

Candidate invariants can also be inferred by generating formulas of some user-provided patterns and testing them against some particular executions of the program [15]. Those formulas that are violated in any of the runs can be rejected, and the soundness of the remaining candidates can be checked afterwards. Invariants can be computed by using iterative backward algorithms [27], starting from the post-condition and computing weakest pre-conditions until a fixpoint is reached (if any). Other approaches [21] have explored the use of quantifier elimination to refine properties obtained using a representation of all execution paths.

The work that is closest to our approach is [13], which presents an algorithm to compute invariants as boolean combinations of linear constraints over integers. The algorithm is similar to ours, and also uses abduction to strengthen candi-

¹ An implicant of a formula ψ is a formula ϕ such that $\phi \models \psi$. It is the dual notion of that of implicants

date invariant so that verification conditions are satisfied. The algorithms differ by the way the verification conditions and abductive hypotheses are proceeded: in our approach the conditions always propagate forward from an invariant to another along execution paths, and we eagerly ensure that all the loop invariants are inductive. Another difference is that we use a completely different technique to perform abductive reasoning: in [13] is based on model construction and quantifier elimination for Presburger arithmetic, whereas our approach uses a generic algorithm, assuming only the existence of a decision procedure for the underlying theory. This permits to generate invariants expressed in theories that are out of the scope of [13].

Contribution. The main contribution is the implementation of a general framework for the generation of loop invariants, connecting the platform Why3 and GPiD. The evaluation demonstrates that the system permits to generate loop invariants for a wide range of theories, though it suffers from a large search space which may induce a large computation time.

2 Verification Conditions

In what follows, we consider formulas in a base logic expressing properties of the memory and assume that such formulas are closed under the usual boolean connectives. These formulas are interpreted modulo some theory \mathcal{T} , where $\models_{\mathcal{T}}$ denotes logical entailment w.r.t. \mathcal{T} . The memory is modified by programs, which are sequences of instructions; they are inductively defined as follows:

$$\begin{aligned} P &= \text{empty} \mid I ; P' \\ I &= \langle \text{base-instruction} \rangle \mid \text{assume } \phi \mid \text{assert } \phi \\ &\mid \text{if } C \text{ then } P_1 \text{ else } P_2 \mid \text{while } C \text{ do } P_1 \{ \phi \} \text{ end} \end{aligned}$$

where P' , P_1 and P_2 are programs, C is a condition on the state of the memory, ϕ is a formula and I is an instruction. Assumptions correspond to formulas that are taken as hypotheses, they are mostly useful to specify pre-conditions. Assertions correspond to formulas that are to be proved. Base instructions are left unspecified, they depend on the target language and application domain; they may include, for instance, assignments and pointer redirection. The formula ϕ in the **while** loop is a *candidate loop invariant*, it is meant to hold every time condition C is tested. In our setting each candidate loop invariant will be set to \top before invoking ILINVA (except when another formula is provided by, *e.g.*, the user), and the program will iteratively update these formulas. We assume that conditions contain no instructions, *i.e.*, that the evaluation of these conditions does not affect the memory. We write $P \sim P'$ if programs P and P' are identical up to the loop candidate invariants.

An example of a program is provided in Figure 1. It uses assignments on integers and usual constructors and functions on lists as base instructions. It contains one loop with candidate invariant \top (Line 3) and one assertion (Line 6).

```

1 let i  $\leftarrow$  1 ;
2 let L  $\leftarrow$  list(1, nil) ;
3 while unknown() do { $\top$ }
4   i  $\leftarrow$  i + 1 ;
5   L  $\leftarrow$  list(i, L) ;
6 assert head(L) = length(L) ;

```

Fig. 1: A simple program on lists

It contains one loop for which we will generate an invariant.

A *location* is a finite sequence of natural numbers. The empty location is denoted by ε and the concatenation of two locations ℓ and ℓ' is denoted by $\ell.\ell'$. If ℓ is a location and S is a set of locations then $\ell.S$ denotes the set $\{\ell.\ell' \mid \ell' \in S\}$. The set of locations in a program P or in an instruction I is inductively defined as follows:

- If P is an empty sequence then $\text{loc}(P) = \{0\}$.
- If $P = I ; P'$ then $\text{loc}(P) = \{0\} \cup 0.\text{loc}(I) \cup \{(i+1).p \mid i \in \mathbb{N}, i.p \in \text{loc}(P')\}$.
- If I is a base instruction or an assumption/assertion, then $\text{loc}(I) = \emptyset$.
- If $I = \text{if } C \text{ then } P_1 \text{ else } P_2$ then $\text{loc}(I) = 1.\text{loc}(P_1) \cup 2.\text{loc}(P_2)$.
- If $I = \text{while } C \text{ do } P_1 \{ \phi \} \text{ end}$ then $\text{loc}(I) = 1.\text{loc}(P_1)$.

For instance, a program $I_1 ; I_2$ where I_1, I_2 denote base instructions has three locations: 0 (beginning of the program), 1 (between I_1 and I_2) and 2 (end of the program). Note that there are no locations within an atomic instruction. The program in Figure 1 has eight locations, namely 0, 1, 2, 2.1.0, 2.1.1, 2.1.2, 3, 4. We denote by $P|_\ell$ the instruction occurring just after location ℓ in P (if any):

- If $P = I ; P'$ then $P|_0 = I$, $P|_{0.\ell} = I|_\ell$ and $P|_{(i+1).\ell} = P'|_{i.\ell}$.
- If $I = \text{if } C \text{ then } P_1 \text{ else } P_2$ then $I|_{1.\ell} = P_1|_\ell$ and $I|_{2.\ell} = P_2|_\ell$.
- If $I = \text{while } C \text{ do } P_1 \{ \phi \} \text{ end}$ then $I|_{1.\ell} = P_1|_\ell$.

Note that $\ell \mapsto P|_\ell$ is a partial function, since locations denoting the end of a sequence do not correspond to an instruction. We denote by $\text{lloc}(P)$ the set of locations ℓ in P such that $P|_\ell$ is a loop and by $\text{loops}(P) = \{P|_\ell \mid \ell \in \text{lloc}(P)\}$ the set of loops occurring in P . For instance, if P denotes the program in Figure 1, then $P|_1$ is **let** $L \leftarrow \text{list}(1, \text{nil})$, and $\text{lloc}(P) = \{2\}$.

We denote by $<$ the usual order on locations: $\ell < \ell'$ iff either there exist numbers i, j and locations ℓ_1, ℓ_2, ℓ_3 such that $\ell = \ell_1.i.\ell_2$, $\ell' = \ell_1.j.\ell_3$ and $i < j$, or there exists a location ℓ'' such that $\ell' = \ell.\ell''$.

We assume the existence of a procedure VCgen that, given a program P , generates a set of *verification conditions* for P . These verification conditions are formulas of the form $\phi \Rightarrow \psi$, each of which is meant to be valid. Given a program P , the set of conditions $\text{VCgen}(P)$ can be decomposed as follows:

1. *Assertion conditions*, which ensure that the assertion formulas hold at the corresponding location in the program. These conditions also include additional properties to prevent memory access errors, *e.g.*, to verify that the index of an array is within the defined valid range of indexes. The set of assertion conditions for program P is denoted by $\text{VCgen}_a(P)$.
2. *Propagation conditions*, ensuring that loop invariants do propagate. Given a loop L occurring at position ℓ in program P , we denote by $\text{VCgen}_{\text{ind}}(P, \ell)$ the set of assertions ensuring that the loop invariant for L propagates.

$$\begin{array}{l}
\hline
wp(\phi, \text{empty}) = \phi \\
wp(\phi, I ; P) = wp(wp(\phi, P), I) \\
wp(\phi, \text{assume } \phi') = \phi' \Rightarrow \phi \\
wp(\phi, \text{assert } \phi') = \phi' \wedge \phi \\
wp(\phi, \text{if } C \text{ then } P_1 \text{ else } P_2) = C \Rightarrow wp(\phi, P_1) \wedge \neg C \Rightarrow wp(\phi, P_2) \\
\hline
wp(\phi, \text{while } C \text{ do } P_1 \{ \psi \} \text{ end}) = \psi \wedge \forall \mathbf{x}. (\psi \Rightarrow wp(\psi, P_1)) \wedge \forall \mathbf{x}. (\psi \wedge \neg C \Rightarrow \phi)
\end{array}$$

The formula in the last line states that the loop invariant holds when the loop is entered, that it propagates and that it entails the formula ϕ . The vector \mathbf{x} denotes the vector of variables occurring in P_1 .

Fig. 2: A Weakest Precondition Calculus

$$\begin{array}{l}
\hline
sp(\phi, \text{empty}) = \phi \\
sp(\phi, I ; P') = sp(sp(\phi, I), P') \\
sp(\phi, \text{assume } \phi') = \phi \wedge \phi' \\
sp(\phi, \text{assert } \phi') = \phi \\
sp(\phi, \text{if } C \text{ then } P_1 \text{ else } P_2) = sp(\phi \wedge C, P_1) \vee sp(\phi \wedge \neg C, P_2) \\
\hline
sp(\phi, \text{while } C \text{ do } P_1 \{ \psi \} \text{ end}) = \psi \wedge \neg C
\end{array}$$

$sp(\phi, P)$ describes the state of the memory after P . The conditions corresponding to loops are approximated by using the provided loop invariants (the corresponding verification conditions are not stated).

Fig. 3: A Strongest Postcondition Calculus

3. *Loop pre-conditions*, ensuring that the loop invariants hold when the corresponding loop is entered. Given a loop L occurring at position ℓ in program P , we denote by $\text{VCgen}_{\text{init}}(P, \ell)$ the set of assertions ensuring that the loop invariant holds before loop L is entered.

Thus, $\text{VCgen}(P) = \text{VCgen}_a(P) \cup \left(\bigcup_{\ell \in \text{lloc}(P)} (\text{VCgen}_{\text{ind}}(P, \ell) \cup \text{VCgen}_{\text{init}}(P, \ell)) \right)$. Such verification conditions are generally defined using standard weakest precondition or strongest post-condition calculi (see, *e.g.*, [12]), where loop invariant are used as under-approximations. Formal definitions are recalled in Figures 2 and 3 (the definition for the basic instructions depends on the application language and is thus omitted). For the sake of readability, we assume, by a slight abuse of notation, that the condition C is also a formula in the base logic.

This permits to define the goal of the paper in a more formal way: our aim is to define an algorithm that, given a program P , constructs a program $P' \sim P$ (*i.e.*, constructs loop invariants for each loop in P) such that $\text{VCgen}(P')$ only contains valid formulas. Note that all the loops and invariants must be handled globally since verification conditions depend on one another.

Algorithm 1: GPID(ϕ, M, A, \mathcal{P})

```
1 if  $M$  unsatisfiable (modulo  $\mathcal{T}$ ) or  $\neg\mathcal{P}(M)$  then
2   return  $\emptyset$ ;
3 if  $M \models \phi$  then
4   return  $\{M\}$ ;
5 let  $\mathbf{m}$  be a model of  $\{\neg\phi\} \cup M$ ;
6 let  $\phi = \text{SIMPLIFY}(\phi, M)$ ;
7 let  $A = \{l \in A \mid M \cup \neg\phi \not\models_{\mathcal{T}} l, M \not\models_{\mathcal{T}} l^c\}$ ;
8 foreach  $l \in A$  such that  $\mathbf{m} \not\models l$  do
9   let  $A_l = \{l' \in A \mid l' < l \wedge \mathbf{m} \models l'\} \cup \{l' \in A \mid l < l'\}$ ;
10  let  $P_l = \text{GPID}(\phi, M \cup \{l\}, A_l, \mathcal{P})$ ;
11 return  $\bigcup_{l \in A} P_l$ ;
```

3 Abduction

As mentioned above, abductive reasoning will be performed by generating implicants. Because it would not be efficient to blindly generate all implicants of a formula, this generation is controlled by fixing the literals that can occur in an implicant. We thus consider a set \mathcal{A} of literals in the considered logic, called the *abducible literals*.

Definition 1. Let ϕ be a formula. An \mathcal{A} -implicant of ϕ (modulo \mathcal{T}) is a conjunction (or set) of literals $l_1 \wedge \dots \wedge l_n$ such that $l_i \in \mathcal{A}$, for all $i \in \llbracket 1 \dots n \rrbracket$ and $l_1 \wedge \dots \wedge l_n \models_{\mathcal{T}} \phi$.

We use the procedure GPID described in [14] to generate \mathcal{A} -implicants. A simplified version of this procedure is presented in Algorithm 1. A call to the procedure GPID(ϕ, M, A, \mathcal{P}) is meant to generate \mathcal{A} -implicants of ϕ that: (i) are of the form $M \cup A'$, for some $A' \subseteq A$; (ii) are as general as possible; and (iii) satisfy property \mathcal{P} . When M itself is not an \mathcal{A} -implicant of ϕ , a subset of relevant literals from A is computed (Line 7), and for each literal in this subset, a recursive call is made to the procedure after augmenting M with this literal and discarding all those that become irrelevant (Lines 9 and 10). In particular, the algorithm is parameterized by an ordering $<$ on abducible literals which is used to ensure that sets of hypotheses are explored in a non-redundant way. The algorithm relies on the existence of a decision procedure for testing satisfiability in \mathcal{T} (Line 1). In practice, this procedure does not need terminating or complete², *e.g.*, it may be called with a timeout (any “unknown” result is handled as “satisfiable”). At Line 8, a model of the formula $\{\neg\phi\} \cup M$ is used to prune the search space, by dismissing some abducible literals. In practice, no such model may be available, either because no model building algorithm exists for the considered theory or because of termination issues. In this case, no such pruning is performed. Property \mathcal{P} denotes an abstract property of sets of literals. It is used to control the

² However, Theorem 2 only holds if the proof procedure is terminating and complete.

form of generated \mathcal{A} -implicants, it is for example possible to force the algorithm to only generate \mathcal{A} -implicants with a fixed maximal size. For Theorem 2 to hold, it is simply required that \mathcal{P} be *closed under subsets*, i.e., that for all sets of abducible literals B and C , $B \subseteq C \wedge \mathcal{P}(C) \Rightarrow \mathcal{P}(B)$.

Compared to [14], details that are irrelevant for the purpose of the present paper are skipped and the procedure has been adapted to generate \mathcal{A} -implicants instead of implicates (implicants and implicates are dual notions).

Theorem 2 ([14]). *The call $\text{GPID}(\phi, \emptyset, \mathcal{A}, \mathcal{P})$ terminates and returns a set of \mathcal{A} -implicants of ϕ satisfying \mathcal{P} . Further, if \mathcal{P} is closed under subsets, then for every \mathcal{A} -implicant I of ϕ satisfying \mathcal{P} , there exists $I' \in \text{GPID}(\phi, \emptyset, \mathcal{A}, \mathcal{P})$ such that $I \models_{\mathcal{T}} I'$.*

This procedure also comes with generic algorithms for pruning redundant \mathcal{A} -implicants i.e., for removing all \mathcal{A} -implicants I such that there exist another \mathcal{A} -implicant I' such that $I \models_{\mathcal{T}} I'$, see [14, Section 4].

4 Generating Loop Invariants

In this section, we present an algorithm for the generation of loop invariants. As explained in Section 2, we distinguish between 3 kinds of verification conditions, which will be handled in different ways: assertion and propagation conditions; and loop pre-conditions. As can be seen from the rules in Figure 2, loop invariants can occur as antecedents in verification conditions, this is typically the case when a loop occurs just before an assertion in some execution path. In such a situation, we say that the considered condition *depends on* loop L . When a condition depends on a loop, a strengthening of the loop invariant of loop L yields a strengthening of the hypotheses of the verification condition, i.e., makes the condition less general (easier to prove).

This principle is used in Algorithm 2, which we briefly describe before going into details. Starting with a program P in which it is assumed that every loop invariant is inductive, the algorithm attempts to recursively generate invariants that make all assertion conditions in P valid. It begins by selecting a non-valid formula ϕ from $\text{VCgen}_a(P)$ and a location $\ell \in \text{lloc}(P)$ such that ϕ depends on ℓ , then generates a set of hypotheses that would make ϕ valid (Line 4). For each such hypothesis ξ , a loop location ℓ' such that $\ell' \leq \ell$ is selected, and a formula ξ' that is a weakest precondition at ℓ' causing ξ to hold at location ℓ is computed (Line 7). This formula is added to the invariant of the loop at location ℓ' (Line 8), so that if this invariant was ψ , the new candidate invariant is $\xi' \wedge \psi$. If ξ' does not hold before entering the loop then ξ is discarded (Line 9); otherwise, the program attempts to update the other loop invariants to ensure that ξ' propagates (Line 10). When this succeeds, a recursive call is made with the updated invariants (Line 12) to handle the other non-valid assertion conditions.

Procedure $\text{ABDUCE}(\phi)$ (invoked Line 4 of Algorithm 2) is described in Algorithm 3. It generates formulas ξ that logically entail ϕ ; it is used to generate the candidate hypotheses for strengthening. It first extracts a set of abducible

Algorithm 2: ILINVA (Program P)

```

1 if all formulas in  $\text{VCgen}_a(\mathbf{P})$  are valid then
2   return P;
3 let  $\phi$  be a non valid formula in  $\text{VCgen}_a(\mathbf{P})$ , depending on a loop at location  $\ell$ ;
4 let  $\Xi \leftarrow \text{ABDUCE}(\phi, \mathbf{P}, \ell)$ ;
5 foreach  $\xi \in \Xi$  do
6   foreach  $\ell' \in \text{lloc}(\mathbf{P})$  such that  $\ell' \leq \ell$  do
7     let  $\xi' \leftarrow bp(\xi, \mathbf{P}, \ell, \ell')$ ;
8     let  $\mathbf{P}_\xi \leftarrow \text{Strengthen}(\mathbf{P}, \ell', \xi')$ ;
9     if  $\text{VCgen}_{\text{init}}(\mathbf{P}_\xi, \ell')$  is valid then
10      let  $\mathbf{P}'_\xi \leftarrow \text{IND}(\mathbf{P}_\xi, \ell')$ ;
11      if  $\mathbf{P}'_\xi \neq \text{fail}$  then
12        let  $\mathbf{P}''_\xi \leftarrow \text{ILINVA}(\mathbf{P}'_\xi)$ ;
13        if  $\mathbf{P}''_\xi \neq \text{fail}$  then
14          return  $\mathbf{P}''_\xi$ ;
15 return fail ;

```

Algorithm 3: ABDUCE(Formula ϕ , Program P, Location ℓ)

```

1 let  $\mathcal{A} \leftarrow \text{GETABDUCIBLES}(\phi)$  ;
2 let  $\mathcal{A} \leftarrow \{l \mid l \in \mathcal{A} \wedge \phi \not\models_{\mathcal{T}} l\}$  ;
3 let  $\Xi \leftarrow \text{GPID}(\phi, \emptyset, \mathcal{A}, \mathcal{P})$  ;
4 let  $\Xi' \leftarrow \{\xi_1 \vee \dots \vee \xi_n \mid n \in \mathbb{N}, \xi_i \in \Xi\}$  ;
5 return  $\Xi'$ 

```

literals \mathcal{A} by collecting variables and symbols from the program and/or from the theory \mathcal{T} and combining them to create literals up to a certain depth (procedure GETABDUCIBLES at Line 1). To avoid any redundancy, this task is actually done in two steps: a set of abducible literals for the entire program is initially constructed (this is done once at the beginning of the search), and depending on the considered program location, a subset of these literals is selected. The abducible literals that are logically entailed by ϕ modulo \mathcal{T} are filtered out (Line 2), and procedure GPID is called to generate \mathcal{A} -implicants of ϕ . Finally, \mathcal{A} -implicants are combined to form disjunctive formulas. Note that another way of generating disjunction of literals would be to add these disjunction in the initial set of abducible literals, but this solution would greatly increase the search space.

Each of the hypotheses ξ generated by ABDUCE(ϕ) is used to strengthen the invariant of a loop occurring at position $\ell' \leq \ell$ (Line 8 in Algorithm 2). The strengthening formula is computed using the Weakest Precondition Calculus on ξ , on a program obtained from P by ignoring all loops between ℓ' and ℓ , since they have corresponding invariants. To this purpose we define a function $bp(\phi, \mathbf{P}, \ell, \ell')$ which, for positions $\ell' \leq \ell$, back-propagates abductive hypotheses from a location ℓ to ℓ' (see Figure 4). This is done by extracting the part of

$$\begin{aligned}
& \text{path}(\mathbf{P}, \ell, \ell) = \text{empty} \\
& \text{path}(\mathbf{P}, \ell, \ell'.(i+1)) = \text{path}(\mathbf{P}, \ell, \ell'.i) \bullet \mathbf{P}|_{\ell'.i} \quad \text{if } \ell \leq \ell'.i \\
& \text{path}(\mathbf{P}, \ell, \ell'.0) = \text{path}(\mathbf{P}, \ell, \ell') \quad \text{if } \ell \leq \ell' \\
& \text{path}(\mathbf{P}, \ell.i.\ell', \ell.(i+1)) = \text{path}(\mathbf{P}, \ell.i.\ell', \ell.i.m) \quad m = \max\{j \mid \ell.i.j \in \text{loc}(\mathbf{P})\} \\
& bp(\phi, \mathbf{P}, \ell, \ell') = wp(\phi, \mathbf{P}') \quad \text{if } \mathbf{P}' = \text{path}(\text{RmLoops}(\mathbf{P}), \ell', \ell) \\
& fp(\phi, \mathbf{P}, \ell, \ell') = sp(\phi, \mathbf{P}') \quad \text{if } \mathbf{P}' = \text{path}(\text{RmLoops}(\mathbf{P}), \ell, \ell')
\end{aligned}$$

$\text{RmLoops}(\mathbf{P})$ denotes the program obtained from \mathbf{P} by removing all **while** instructions and \bullet denotes the concatenation operator on programs.

Fig. 4: Backward and Forward Propagation of Abductive Hypotheses

Algorithm 4: IND (Program \mathbf{P} , Location ℓ)

```

1 if all formulas in  $\text{VCgen}_{\text{ind}}(\mathbf{P}, \ell)$  are valid then
2   return  $\mathbf{P}$ ;
3 let  $\phi$  be a non-valid formula in  $\text{VCgen}_{\text{ind}}(\mathbf{P}, \ell)$  ;
4 let  $\Xi \leftarrow \text{ABDUCE}(\phi, \mathbf{P}, \ell)$ ;
5 foreach  $\xi \in \Xi$  do
6   foreach  $\ell' \in \text{loc}(\mathbf{P})$  such that  $\ell$  is a prefix of  $\ell'$  (with possibly  $\ell = \ell'$ ) do
7     let  $\xi' \leftarrow fp(\xi, \mathbf{P}, \ell, \ell')$ ;
8     let  $\mathbf{P}'_{\xi} \leftarrow \text{Strengthen}(\mathbf{P}, \ell', \xi')$  ;
9     if  $\text{VCgen}_{\text{init}}(\mathbf{P}'_{\xi}, \ell')$  is valid then
10      let  $\mathbf{P}''_{\xi} \leftarrow \text{IND}(\mathbf{P}'_{\xi}, \ell)$  ;
11      if  $\mathbf{P}''_{\xi} \neq \text{fail}$  then
12        return  $\mathbf{P}''_{\xi}$ ;
13 return fail ;

```

the code $\text{path}(\mathbf{P}, \ell', \ell)$ between the locations ℓ' and ℓ while ignoring loops, and computing the weakest precondition corresponding to this part of the code and the formula ϕ .

The addition of hypothesis ξ' to the invariant of the loop at position ℓ' ensures that the considered assertion ϕ holds, but it is necessary to ensure that this strengthened invariant is still inductive. This is done as follows. Line 9 of Algorithm 2 filters away all candidates for which the precondition before entering the loop is no longer valid, and Algorithm 4 ensures that the candidate still propagates. This algorithm behaves similarly to Algorithm 2 (testing the verification conditions in $\text{VCgen}_{\text{ind}}(\mathbf{P}, \ell)$ instead of those in $\text{VCgen}_{\text{a}}(\mathbf{P})$), except that it strengthens the invariants that correspond either to the considered loop, or to other loops occurring within it (in the case of nested loops). Note that in this case, properties must be propagated forward, from location ℓ to the actual location of the strengthened invariant, using a Strongest Postcondition Calculus (Function $fp(\phi, \mathbf{P}, \ell, \ell')$ in Figure 4). This technique avoids considering hypotheses that do not propagate.

When applied on the program in Figure 1, ILINVA first sets the initial invariant of the loop to \top and considers the assertion $\phi : \mathbf{head}(L) = \mathbf{length}(L)$. As the entailment $\top \models \phi$ does not hold, it will call GPID to get an implicant of $\top \Rightarrow \phi$. Assume that GPID returns the (trivial) solution ϕ . As ϕ indeed holds when the loop is entered³, ILINVA will add ϕ to the invariant of the loop and call IND. Since ϕ does not propagate IND will further strengthen the invariant, yielding, *e.g.*, the correct solution: $\phi \wedge i = \mathbf{head}(L)$.

The efficiency of Algorithm 2 crucially depends on the order in which candidate hypotheses are processed at Line 5 for the strengthening operation. The heuristic used in our current implementation is to try the simplest hypotheses with the highest priority. Abducible atoms are therefore ordered as follows: first boolean variables, then equations between variables of the same sort, then applications of predicate symbols to variables (of the appropriate sorts) and finally deep literals involving function symbols (up to a certain depth). In every case, negative literals are also considered, with the same priority as the corresponding atom. Similarly, unit \mathcal{A} -implicants are tested before non-unit ones, and single \mathcal{A} -implicants before disjunctions of \mathcal{A} -implicants. In the iteration on line 6 of Algorithm 2, the loops that are closest to the considered assertions are considered first. Due to the number of loops involved, numerous parameters are used to control the application of the procedures, by fixing limits on the number of abducible literals that may be considered and on the maximal size of \mathcal{A} -implicants. When a call to ILINVA fails, these parameters are increased, using an iterative deepening search strategy. The parameter controlling the maximal number of \mathcal{A} -implicants in the disjunctions (currently either 1 or 2) is fixed outside of the loop as it has a strong impact on the computation cost.

The following theorem states the main properties of the algorithm.

Theorem 3. *Let P be a program such that $\text{VCgen}_{\text{ind}}(P, \ell)$ and $\text{VCgen}_{\text{init}}(P, \ell)$ are valid for all $\ell \in \text{lloc}(P)$. If $\text{ILINVA}(P)$ terminates and returns a program P' other than **fail**, then $P \sim P'$ and $\text{VCgen}(P')$ is valid modulo \mathcal{T} . Furthermore, if the considered set of abducible literals is finite (i.e., if there exists a finite set \mathcal{A} such that $\text{GETABDUCIBLES}(\phi) \subseteq \mathcal{A}$ for all formulas ϕ), then $\text{ILINVA}(P)$ terminates.*

Proof. The proof is provided in the extended version⁴.

5 Implementation

5.1 Overview

The ILINVA algorithm described in Section 4 has been implemented by connecting Why3 with GPID. A workflow graph of this implementation is detailed in the

³ This can be checked by computing the weakest precondition of ϕ w.r.t. Lines 1, 2. The obtained formula is $\mathbf{head}(\mathbf{list}(1, \mathbf{nil})) = \mathbf{length}(\mathbf{list}(1, \mathbf{nil}))$ which is equivalent to \top (w.r.t. the usual definitions of **list** and **head**).

⁴ <https://arxiv.org/abs/1906.11033>

extended paper. Note that both systems themselves call external SMT solvers to check the satisfiability of formulas. In particular, the GPiD toolbox is easy to plug to any SMTLIB2-compliant SMT solver. The framework is actually generic, in the sense that it could be plugged with other systems, both to generate and verify proof obligations and to strengthen loop invariants. It is also independent of the constructions used for defining the language: other constructions (*e.g.*, `for` loops) can be considered, provided they are handled by the program verification framework.

Given an input program written in WHYML, `Why3` generates a verification condition the validity of which will ensure that all the asserted properties are verified (including additional conditions related to, *e.g.*, memory safety). This initial verification condition is split by `Why3` into several subtasks. These conditions are enriched with all appropriate information (*e.g.*, theories, axioms, ...) and sent to external SMT solvers to check satisfiability. The conditions we are interested in are those linked to the proofs of the program assertions, as well as those ensuring that the candidate loop invariants are inductive. In our implementation, `Why3` is taken as a black box, and we merely recover the files that are passed from `Why3` to the SMT solvers, together with additional configuration data for the solvers we can extract from `Why3`. If the proof obligation fails, then we relate the file to the corresponding assertion in the WHYML program and extract the set of abducible literals as explained in Section 4, restricting ourselves to symbols corresponding to WHYML variables, functions and predicates. We then tune the SMTLIB2 file to adapt it for computations by GPiD and invoke GPiD with the same SMT-solver as the one used by `Why3` to check satisfiability, as the problem is expected to be optimized/configured for it. We also configure GPiD to skip the exploration of subtrees that will produce candidate invariants that do not satisfy the loop preconditions. GPiD returns a stream of solutions to the abductive reasoning problem. We then backward-translate the formulas into the WHYML formalism and use them to strengthen loop invariants. For efficiency, the systems run in parallel: the generation of abductive hypotheses (by GPiD, via the procedure `ABDUCE`) and their processing in WHYML (via `ILINVA`) is organized as a pipe-line, where new abduction solutions are computed during the processing of the first ones.

To bridge `ILINVA` and `Why3`, we had to devise an interface, which is able to analyze WHYML programs and to identify loop locations and the corresponding invariants. It invokes `Why3` to generate and prove the associated verification tasks, and it recovers the failed ones. The library also includes tools to extract and modify loop invariants, to extract variables and reference variables in WHYML files, as well as types, predicates and functions, and wrappers to call the `Why3` executable and external tools, and to extract the files sent by WHYML to SMT-solvers.

5.2 Distribution

The ABDULOT framework is available on GITHUB [7]. It contains an revamped interface to the GPiD libraries and algorithm, a generic library of the `ILINVA`

algorithm automatically plugged with GPiD, the code interface for Why3 and the related executables. GPiD interfaces and related executables are generated for CVC4, Z3 and ALTERGO⁵ via their SMTLIB2 interface. Note that the SMT solvers are not provided by our framework, they must be installed separately (all versions with an SMTLIB2-compatible interface are supported). Additional interfaces and executables can be produced using C++ libraries for MINISAT, CVC4 and Z3 if their supported version is available⁶.

The framework also provides libraries and toolbox executables to work with abducible files, C++ libraries to handle WHYML files, helpers for the generation of abducible literals out of SMTLIB2 files, and an extensive lisp parser. It also includes a documentation, which explains in particular how to extend it to other solvers and program verification framework. All the tools can be compiled using any C++ 11 compliant compiler. The whole list of dependencies is available in the documentation, as well as a dependency graph for the different parts of the framework.

6 Experiments

We use benchmarks collected from several sources [13,4,5,6,1,2,3] (see also [7] for a more detailed view of the benchmark sources), with additional examples corresponding to standard algorithms for division and exponentiation (involving lists, arrays, and non linear arithmetic). Some of these benchmarks have been translated⁷ from C or JAVA into WHYML. In all cases, the initial invariant associated with each loop is \top . We used Z3 for the benchmarks containing real arithmetic, ALTERGO for lists and arrays and CVC4 in all the other cases. All examples are embedded with the source of the `Ilinva` tool.

6.1 Results

We ran `Ilinva` on each example, first without disjunctive invariants (*i.e.*, taking $n = 1$ in Procedure ABDUCE) then with disjunctions of size 2. The results are reported in Figure 5. For each example, we report whether our tool was able generate invariants allowing Why3 to check the soundness of all program assertions before the timeout, in which case we also report the time `Ilinva` took to do so (columns T(C) when generating conjunctions only and T(D) when generating implicants containing disjunctions). We also report the number of candidate invariants that have been tried (columns C(D) and C(D)) and the number of abducible literals that were sent to the GPiD algorithm (column Abd). Note that the number of candidate invariants does not correspond to the

⁵ Those are the three solvers the Why3 documentation recommends to work with as an initial setup. (see also <http://why3.lri.fr/@ExternalProvers>.)

⁶ The ALTERGO interface provided by the tool uses an SMTLIB2 interface that is under heavy development and that, in practice, does not work well with the examples we send it.

⁷ The translation was done by hand.

number of SMT calls that are made by the system: those made by GPiD to generate these candidates are not taken into account. The timeout is set to 20 min. For some of the examples that we deemed interesting, we allowed the algorithm to run longer. We report those cases by putting the results between parentheses. Light gray cells indicate that the program terminates before the timeout without returning any solution, and dark gray cells indicate that the timeout was reached. Empty cells mean that the tool could not generate any candidate invariant. The last column of both tables report the time `Why3` takes to prove all the assertions of an example when correct invariants are provided.

The tests were performed on a computer powered by a dual-core Intel i5 processor running at 1.3GHz with 4 GB of RAM, under macOS 10.14.3. We used `Why3` version 1.2.0 and the SMT solvers `ALTERGO` (version 2.2.0), `CVC4` (prerelease of version 1.7) and `Z3` (version 4.7.1).

An essential point concerns the handling of local solver timeouts. Indeed, most calls to the SMT solver in the abductive reasoning procedure will involve satisfiable formulas, and the solvers usually take a lot of time to terminate on such formulas (or in the worst case will not terminate at all if the theory is not decidable, *e.g.*, for problems involving first-order axioms). We thus need to set a timeout after which a call will be considered as satisfiable (see Section 3). Obviously, we neither want this timeout to be too high as it can significantly increase computation time, nor too low, since it could make us miss solutions. We decided to set this timeout to 1 second, independently of the solver used, after measuring the computation time of the `Why3` verification conditions already satisfied (for which the solver returns `unsat`) across all benchmarks. We worked under the assumption that the computation time required to prove the other verification conditions when possible would be roughly similar.

6.2 Discussion

As can be observed, `ILINVA` is able to generate solutions for a wide range of theories, although the execution time is usually high. The number of invariant candidates is relatively high, which has a major impact on the efficiency and scalability of the approach.

When applied to examples involving arithmetic invariants, the program is rather slow, compared to the approach based on minimization and quantifier elimination [13]. This is not surprising, since it is very unlikely that a purely generic approach based on a model-based tree exploration algorithm involving many calls to an SMT solver can possibly compete with a more specific procedure exploiting particular properties of the considered theory. We also wish to emphasize that the fact that our framework is based on an external program verification system (which itself calls external solvers) involves a significant overcost compared to a more integrated approach: for instance, for the `Oxy` examples (taken from [13]), the time used by `Why3` to check the verification conditions once the correct invariants have been generated is often greater than the total time reported in [13] for computing the invariants and checking all properties.

	Abd	T(C)	C(C)	T(D)	C(D)	Why3
001	36	9.68	7	11.89	10	0.26
002	536	3'18.9	66		1126	0.45
004	108	50.47	32	2'31.4	156	0.26
005	266	1'9.07	5	1'3.2	5	0.31
006	390	6'13.6	56	18'5.1	552	0.72
007	594	1'50.1	13	15'40.6	355	0.38
008	210	2'35.5	61	9'35.8	528	0.42
009	390		0		0	0.56
010	90	1'39.54	65	12'56.9	0	0.35
011	180	2'17.7	63		942	0.26
012	782		0		0	0.53
013	296	2'4.5	0		1621	0.30
014	270		0		0	0.34
015	36	32.53	21		888	0.27
016	60	12.54	8	29.72	32	0.26
017	36	40.88	26	2'42.5	134	0.33
018	38	58.49	38	6'53.3	0	0.30
019	60	1'59.5	111		1620	0.31
020	546		380		870	0.49
021	90	0.76	0	0.76	0	0.38
022	270	2'10.1	20	2'11.9	20	0.48
023	36	4.6	5	4.7	5	0.28
025	60	1'23.4	20	2'38.4	44	0.39
026	396	6'23.2	21	7'13.9	66	0.83
028		2'3.9	137	16'22.8	1331	0.31
029	61776		0		0	0.65
030	36	31.43	26	41.66	45	0.26
031	67050		0		0	0.49
032	40	0.865	0	0.833	0	0.50
033	90	1'11.3	12	1'19.9	21	0.45
034	6768	0.798	0	0.79	0	0.44
035	18	18.42	25	2'7.9	200	0.26
036	61778		0		0	1.09
037	36	0.752	0	0.769	0	0.34
038	630		444	3'54.4	0	0.48
039	546		1581		1840	0.40
040	272		0		0	0.84
041			0		0	0.37
042	271	1'50.4	25		605	1.12
043	60	4.27	2	3.67	2	0.29
044		22.481	13	5'7.8	290	0.35
045			0		0	1.50
046			513		813	0.61

	Abd	T(C)	C(C)	T(D)	C(D)	Why3
509	130	(1h50')	(95)		0	0.66
534	172k		8		0	0.62
H04	120	2'54.8	223		1383	0.31
H05	1260		0		0	0.37
list0	60	6'30.4	77		1722	0.40
list1	20	40.82	3	3'26.2	385	0.47
list2	720		40		0	0.40
list3	126	3'35.1	11		930	0.44
list4	816		18		0	
list5	468		22		0	0.44
array0			0		0	0.72
array1			0		0	0.50
array2			0		0	0.50
array3			0		0	0.82
expo0	171	(6h36')	(9)		0	0.40
expo1	2130		0		0	
square	705		62		148	
real0	965		81		213	0.55
real1	965		73		115	0.55
real2	240		9		2	0.40
real00	36	4'9.6	25	5'32.18	40	0.47
realS	66	1'5.3	5	1'0.1	5	0.33
real3	17460		0		0	
BM	1260	3'15.2	74		33	3.35
Scmp			0		0	0.83
Dmd	42		6		0	1'44.9
B00	639k		0		0	0.76
DIV0	560	3'58	35		534	0.83
DIV1	310	14.6	19	14.6	19	0.44
DIVE	42250		0		0	

Fig. 5: Experimental Results

Of course, our choice also has clear advantages in terms of genericity, generality and evolvability.

When applied to theories that are harder for SMT solvers, the algorithm can still generate satisfying invariants. However, due to the high number of candidates it tries, combined with the heavy time cost of a verification (which can be several seconds), it may take some time to do so.

The number of abducible literals has a strong impact on the efficiency of the process, leading to timeouts when the considered program contains many variables or function/predicate symbols. It can be observed that the abduction depth is rather low in all examples (1 or 2).

Our prototype has some technical limitations that have a significant impact on the time cost of the execution. For instance, we use SMTLIB2 files for communication between GPID and CVC4 or Z3, instead of using the available APIs. We went back to this solution, which is clearly not optimal for performance, because we experienced many problems coping with the numerous changes in the specifications when updating the solvers to always use the latest versions. The fact that `Why3` is taken as a black box also yields some time consumption, first in the (backward and forward) translations (*e.g.*, to associate program variables to their logical counterparts), but also in the verification tasks, which have to be rechecked from the beginning each time an invariant is updated. Our aim in the present paper was not to devise an efficient system, but rather to assess the feasibility and usefulness of this approach. Still, the cost of the numerous calls to the SMT solvers and the size of the search tree of the abduction procedure remain the bottleneck of the approach, especially for hard theories (*e.g.*, non-linear arithmetics) for which most calls with satisfiable formulas yield to a local timeout (see Section 6.1).

7 Conclusion and Future Work

By combining our generic system GPID for abductive reasoning modulo theories with the `Why3` platform to generate verification conditions, we obtained a tool to check properties of `WHYML` programs, which is able to compute loop invariants in a purely automated way.

The main drawback of our approach is that the set of possible abducible literals is large, yielding a huge search space, especially if disjunctions of \mathcal{A} -implicants are considered. Therefore, we believe that our system in its current state is mainly useful when used in an interactive way. For instance, the user could provide the properties of interest for some of the loops and let the system automatically compute suitable invariants by combining these properties, or the program could rely on the user to choose between different solutions to the abduction problem before applying the strengthening. Beside, it is also useful for dealing with theories for which no specific abductive reasoning procedure exists, especially for reasoning in the presence of user-defined symbols or axioms.

In the future, we will focus on the definition of suitable heuristics for automatically selecting abducible literals and ordering them, to reduce the search

space and avoid backtracking. The number of occurrences of symbols should be taken into account, as well as properties propagating from previous invariant strengthening. A promising approach is to use dynamic program analysis tools to select relevant abducibles. It would also be interesting to adapt the GPiD algorithm to explore the search space width-first, to ensure that simplest solutions are always generated first. Another option is to give ILINVA a more precise control on the GPiD algorithm, *e.g.*, to explore some branches more deeply, based on information related to the verification conditions. GPiD could also be tuned to generate disjunctions of solutions in a more efficient way.

From a more technical point of view, a tighter integration with the Why3 platform would certainly be beneficial, as explained in Section 6.2. The framework could be extended to handle procedures and functions (with pre- and -post conditions).

A current limitation of our tool is that it cannot handle problems in which Why3 relies on a combination of different solvers to check the desired properties. In this case, ILINVA cannot generate the invariants, as the same SMT solver is used for each abduction problem (trying all solvers in parallel on every problem would be possible in theory but this would greatly increase the search space). This problem could be overcome by using heuristic approaches to select the most suited solver for a given problem.

From a theoretical point of view, it would be interesting to investigate the completeness of our approach. It is clear that no general completeness result possibly holds, due to usual theoretical limitations, however, if we assume that a program $P' \sim P$ such that $\text{VCgen}(P')$ is valid exists, does the call $\text{ILINVA}(P)$ always succeed? This of course would require that the invariants in P' can be constructed from abducibles occurring in the set returned by the procedure `GETABDUCIBLES`.

References

1. <http://toccata.lri.fr/gallery/>.
2. <http://pauillac.inria.fr/~levy/why3/sorting/>.
3. <https://www.lri.fr/~sboldo/research.html>.
4. Invgen tool. <http://pub.ist.ac.at/agupta/invgen/>.
5. Neclabs necla verification benchmarks. <http://www.nec-labs.com/research/system/systemsSAV-website/benchmarks.php>.
6. Satconv benchmarks.
7. Abdulot framework/GPiD-Ilinva tool suite. <https://github.com/sellamiy/GPiD-Framework>.
8. D. Beyer, T. A. Henzinger, R. Majumdar, and A. Rybalchenko. Invariant synthesis for combined theories. In *Verification, Model Checking, and Abstract Interpretation, 8th International Conference, VMCAI 2007, Nice, Proceedings*, 2007.
9. A. R. Bradley. IC3 and beyond: Incremental, inductive verification. In *Computer Aided Verification - 24th International Conference, CAV 2012, Berkeley, CA, USA, July 7-13, 2012 Proceedings*, page 4, 2012.
10. A. R. Bradley and Z. Manna. Property-directed incremental invariant generation. *Formal Asp. Comput.*, 20(4-5):379–405, 2008.

11. P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Proceedings of the 5th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, POPL '78, New York, 1978. ACM.
12. E. W. Dijkstra. *A Discipline of Programming*. Prentice Hall PTR, Upper Saddle River, NJ, USA, 1st edition, 1997.
13. I. Dillig, T. Dillig, B. Li, and K. L. McMillan. Inductive invariant generation via abductive inference. In A. L. Hosking, P. T. Eugster, and C. V. Lopes, editors, *Proceedings of OOPSLA 2013, Indianapolis*, pages 443–456. ACM, 2013.
14. M. Echenim, N. Peltier, and Y. Sellami. A generic framework for implicate generation modulo theories. In D. Galmiche, S. Schulz, and R. Sebastiani, editors, *IJCAR 2018, Oxford*, volume 10900 of *LNCS*, pages 279–294. Springer, 2018.
15. M. D. Ernst, J. Cockrell, W. G. Griswold, and D. Notkin. Dynamically discovering likely program invariants to support program evolution. In *Proceedings of the 21st International Conference on Software Engineering*, ICSE '99, pages 213–224, New York, NY, USA, 1999. ACM.
16. J.-C. Filliâtre and A. Paskevich. Why3 — where programs meet provers. In M. Felleisen and P. Gardner, editors, *Proceedings of the 22nd European Symposium on Programming*, volume 7792 of *Lecture Notes in Computer Science*, pages 125–128. Springer, Mar. 2013.
17. C. Flanagan and K. R. M. Leino. Houdini, an annotation assistant for esc/java. In J. N. Oliveira and P. Zave, editors, *FME 2001: Formal Methods for Increasing Software Productivity*, pages 500–517, Berlin, Heidelberg, 2001. Springer Berlin Heidelberg.
18. S. Ghilardi and S. Ranise. Backward reachability of array-based systems by SMT solving: Termination and invariant synthesis. *Logical Methods in Computer Science*, 6(4), 2010.
19. T. A. Henzinger, R. Jhala, R. Majumdar, and G. Sutre. Software verification with blast. In T. Ball and S. K. Rajamani, editors, *Model Checking Software*, pages 235–239, Berlin, Heidelberg, 2003. Springer Berlin Heidelberg.
20. C. A. R. Hoare. An axiomatic basis for computer programming. *Commun. ACM*, 12(10):576–580, Oct. 1969.
21. D. Kapur. A quantifier-elimination based heuristic for automatically generating inductive assertions for programs. *J. Systems Science & Complexity*, 19, 2006.
22. A. Karbyshev, N. Bjørner, S. Itzhaky, N. Rinetzky, and S. Shoham. Property-directed inference of universal invariants or proving their absence. *J. ACM*, 64(1):7:1–7:33, 2017.
23. L. Kovács and A. Voronkov. Finding loop invariants for programs over arrays using a theorem prover. In *Fundamental Approaches to Software Engineering, 12th International Conference, FASE 2009, Held as Part of the Joint European Conferences on Theory and Practice of Software, ETAPS 2009, York, UK, March 22-29, 2009. Proceedings*, pages 470–485, 2009.
24. L. Kovács and A. Voronkov. Interpolation and symbol elimination. In *Automated Deduction - CADE-22, 22nd International Conference on Automated Deduction, Montreal, Canada, August 2-7, 2009. Proceedings*, pages 199–213, 2009.
25. A. Miné. The octagon abstract domain. *Higher Order Symbol. Comput.*, 19, 2006.
26. O. Padon, N. Immerman, S. Shoham, A. Karbyshev, and M. Sagiv. Decidability of inferring inductive invariants. In *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 217–231, 2016.
27. N. Suzuki and K. Ishihata. Implementation of an array bound checker. 1977.