

An iterative approach to precondition inference using constrained Horn clauses

Kafle, Bishoksan; Gallagher, John Patrick; Gange, Graeme; Schachte, Peter; Søndergaard, Harald; Stuckey, Peter J.

Published in:
Theory and Practice of Logic Programming

DOI:
[10.1017/S1471068418000091](https://doi.org/10.1017/S1471068418000091)

Publication date:
2018

Document Version
Early version, also known as pre-print

Citation for published version (APA):
Kafle, B., Gallagher, J. P., Gange, G., Schachte, P., Søndergaard, H., & Stuckey, P. J. (2018). An iterative approach to precondition inference using constrained Horn clauses. *Theory and Practice of Logic Programming*, 18(3-4), 553-570. <https://doi.org/10.1017/S1471068418000091>

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain.
- You may freely distribute the URL identifying the publication in the public portal.

Take down policy

If you believe that this document breaches copyright please contact rucforsk@ruc.dk providing details, and we will remove access to the work immediately and investigate your claim.

An iterative approach to precondition inference using constrained Horn clauses

Bishoksan Kafle

The University of Melbourne

John P. Gallagher

Roskilde University and IMDEA Software Institute

Graeme Gange

Monash University

Peter Schachte, Harald Søndergaard, Peter J. Stuckey

The University of Melbourne

submitted 1 January 2003; revised 1 January 2003; accepted 1 January 2003

Abstract

We present a method for automatic inference of conditions on the initial states of a program that guarantee that the safety assertions in the program are not violated. Constrained Horn clauses (CHCs) are used to model the program and assertions in a uniform way, and we use standard abstract interpretations to derive an over-approximation of the set of *unsafe* initial states. The precondition then is the constraint corresponding to the complement of that set, under-approximating the set of *safe* initial states. This idea of complementation is not new, but previous attempts to exploit it have suffered from loss of precision. Here we develop an iterative specialisation algorithm to give more precise, and in some cases optimal safety conditions. The algorithm combines existing transformations, namely constraint specialisation, partial evaluation and a trace elimination transformation. The last two of these transformations perform polyvariant specialisation, leading to disjunctive constraints which improve precision. The algorithm is implemented and tested on a benchmark suite of programs from the literature in precondition inference and software verification competitions.

Under consideration for acceptance in TPLP.

KEYWORDS: Precondition inference, backwards analysis, abstract interpretation, refinement, program specialisation, program transformation.

1 Introduction

Given a program with properties required to hold at specific program points, *precondition analysis* derives the conditions on the initial states ensuring that the properties hold. This has important applications in program verification, symbolic execution, program understanding and debugging. While forward abstract interpretation approximates the set of reachable states of a program, *backward* abstract

interpretation approximates the set of states that can reach some target state. Both forward and backward analyses may produce over- or under-approximations, and forward and backward analysis may profitably be combined (Cousot and Cousot 1992; Cousot et al. 2011; Bakhirkin and Monniaux 2017).

Most approaches that apply backward analysis, possibly in conjunction with forward analysis, use over-approximations, and as a result derive *necessary* pre-conditions. Less attention has been given to under-approximating backwards analyses, with the goal of finding *sufficient* pre-conditions. However, it is natural to try to derive guarantees of safe behaviour of a program. Often we would like to know which initial states *must* be safe, in the sense that no computation starting from such a state can possibly reach a specified error state, that is, we desire to find (non-trivial) sufficient conditions for safety.

If analysis uses an abstract domain which is *complemented*, duality enables sufficient conditions to be derived from necessary conditions and *vice versa*. However, complemented abstract domains are very rare, and approximation of a complement tends to introduce considerable lack of precision. The under-approximating backward abstract interpretation of Howe et al. (2004) utilises the fact that the abstract domain *Pos* is pseudo-complemented (Marriott and Søndergaard 1993), but pseudo-complementation too is very rare. Moy (2008) presents a method for deriving sufficient preconditions (for use with a theorem prover), employing weakest-precondition reasoning and forward abstract interpretation to attempt to generalise conditions at loop heads. Bakhirkin et al. (2014) observe that there may be an advantage in generalising an abstract complement operation to (abstract) logical subtraction, as this can improve opportunities to find a tighter approximation of a set of states.

Miné (2012a) infers sufficient conditions for safety, not by instantiating a generic mechanism for complementation, but by designing all required purpose-built backward transfer functions. He does this for three numeric abstract domains: intervals, octagons and convex polyhedra—a substantial effort, as the purpose-built operations, including widening, can be rather intricate.

We share Miné’s goal but use program transformation and over-approximating abstract interpretation over a Horn clause program representation. This allows us to apply a range of established tools and techniques beyond abstract interpretation, including query-answer transformation, partial evaluation and abstraction refinement. We offer an iterative approach that successively specialises a program. The approach of iteratively specialising a program represented as Horn clauses has also been pursued by De Angelis et al. (2014) in order to verify program properties. Their techniques also incorporate forward and backward propagation of constraints, but rather than explicitly using abstract interpretation, their specialisation algorithm involves a special constraint generalisation method.

We shall use the example in Figure 1 to demonstrate our approach. The left side shows a C program fragment, and the right its constrained Horn clause (CHC) representation. CHCs can be obtained from an imperative program (containing assertions) using various approaches (Peralta et al. 1998; Grebenschikov et al. 2012; Gurfinkel et al. 2015; De Angelis et al. 2017). The A set of CHCs is not necessarily intended as an executable logic program; in Figure 1 the predicates

<pre> int a, b; if (a ≤ 100) a = 100 - a; else a = a - 100; while (a ≥ 1) {a = a - 1; b = b - 2;} assert(b ≠ 0); </pre>	<pre> c1. init(A, B) ← true. c2. if(A, B) ← A₀ ≤ 100, A = 100 - A₀, init(A₀, B). c3. if(A, B) ← A₀ ≥ 101, A = A₀ - 100, init(A₀, B). c4. while(A, B) ← if(A, B). c5. while(A, B) ← A₀ ≥ 1, A = A₀ - 1, B = B₀ - 2, while(A₀, B₀). c6. false ← A ≤ 0, B = 0, while(A, B). </pre>
--	--

Fig. 1: Running example: (left) original program, (right) translation to CHCs

capture the reachable states of the computation. For example, `while(1, 0)` is true if the `while` statement is reached with $a = 1$ and $b = 0$. The predicate `false` represents an error state. Henceforth whenever we refer to a program, we refer to its CHC version.

For the given program, we want to ensure that b is non-zero after the loop. The goal is to derive initial conditions on a and b , sufficient to ensure that the assertion is never violated. The practical use of the conditions is to reject unsafe initial states before running the program. We note that the assertion will *not* be violated provided the following three conditions are met: (i) if $a = 100$ then $b \neq 0$, (ii) if $a < 100$ then $2a \neq 200 - b$ and (iii) if $a > 100$ then $2a \neq 200 + b$. The conjunction of these three conditions, or equivalently $b \neq |2a - 200|$, ensures that the assertion is never violated. Automating the required reasoning is challenging because: (i) the desired result is a disjunctive constraint over expressions that need an expressive domain; (ii) the disjuncts cannot be represented as intervals, octagons or difference bound matrices (Miné 2006); (iii) information has to be propagated forwards and backwards because we lose information on b and a in the forward and in the backward direction respectively. In what follows, we show how to derive the conditions automatically.

The key contribution of this paper is a framework for deriving sufficient preconditions without a need to calculate weakest preconditions or rely on abstract domains with special properties or intricate transfer functions. This is achieved through a combination of program transformation and abstract interpretation, with the derived preconditions being successively refined through iterated transformation.

After Section 2’s preliminaries, we discuss, in Section 3.1, the required transformation techniques. Section 3.2 gives iterative refinement algorithms that derive successively better (weaker) preconditions. Section 4 is an account of experimental evaluation, demonstrating practical feasibility of the technique. Section 5 concludes.

2 Preliminaries

An atomic formula, or simply *atom*, is a formula $p(\mathbf{x})$ where p is a predicate symbol and \mathbf{x} a tuple of arguments. A constrained Horn clause (CHC) is a first-order predicate logic formula of the form $\forall \mathbf{x}_0 \dots \mathbf{x}_k (p_1(\mathbf{x}_1) \wedge \dots \wedge p_k(\mathbf{x}_k) \wedge \phi \rightarrow p_0(\mathbf{x}_0))$, where ϕ is a finite conjunction of quantifier-free *constraints* on variables \mathbf{x}_i with respect to some constraint theory \mathbb{T} , $p_i(\mathbf{x}_i)$ are atoms, $p_0(\mathbf{x}_0)$ is the *head* of the clause and $p_1(\mathbf{x}_1) \wedge \dots \wedge p_k(\mathbf{x}_k) \wedge \phi$ is the *body*. Following the conventions of Constraint Logic

Programming (CLP), such a clause is written as $p_0(\mathbf{x}_0) \leftarrow \phi, p_1(\mathbf{x}_1), \dots, p_k(\mathbf{x}_k)$. For concrete examples of CHCs we use Prolog-like syntax and typewriter font, with capital letters for variable names and linear arithmetic constraints built with predicates $\leq, \geq, <, >, =$.

An *Integrity constraint* is a special kind of clause whose head is the predicate `false`. A *constrained fact* is a clause of the form $p_0(\mathbf{x}_0) \leftarrow \phi$. A set of CHCs is also called a program.

Figure 1 (right) contains an example of a set of constrained Horn clauses. The first five clauses define the behaviour of the program in Figure 1 (left) and the last clause represents a property of the program (that the variable B is non-zero after executing the program) expressed as an integrity constraint.

CHC semantics. The semantics of CHCs is obtained using standard concepts from predicate logic semantics. An *interpretation* assigns to each predicate a relation over the domain of the constraint theory \mathbb{T} . The predicate `false` is always interpreted as *false*. We assume that \mathbb{T} is equipped with a decision procedure and a projection operator, and that it is closed under negation. We use notation $\phi|_V$ to represent the constraint formulae ϕ projected onto variables V .

An interpretation *satisfies* a set of formulas if each formula in the set evaluates to *true* in the interpretation in the standard way. In particular, a *model* of a set of CHCs is an interpretation in which each clause evaluates to *true*. A set of CHCs P is *consistent* if and only if it has a model. Otherwise it is *inconsistent*.

When modelling safety properties of systems using CHCs, the consistency of a set of CHCs corresponds to *safety* of the system. Thus we also refer to CHCs as being *safe* or *unsafe* when they are consistent or inconsistent respectively.

AND-trees and trace trees. Derivations for CHCs are represented by AND-trees. The following definitions are adapted from Gallagher and Lafave (1996).

An *AND-tree* for a set of CHCs is a tree whose nodes are labelled as follows.

1. each non-leaf node corresponds to a clause (with variables suitably renamed) of the form $A \leftarrow \phi, A_1, \dots, A_k$ and is labelled by the atom A and ϕ , and has children labelled by A_1, \dots, A_k ;
2. each leaf node corresponds to a clause of the form $A \leftarrow \phi$ (with variables suitably renamed) and is labelled by the atom A and ϕ ; and
3. each node is labelled with the clause identifier of the corresponding clause.

Of particular interest are AND-trees with their roots labelled by the atom *false*; these are called counterexamples. A *trace tree* is the result of removing all node labels from an AND-tree apart from the clause identifiers. Given an AND-tree t , $\text{constr}(t)$ represents the conjunction of the constraints in its node labels. The tree t is *feasible* if and only $\text{constr}(t)$ is satisfiable over \mathbb{T} . We also represent a conjunction of constraints as a set of constraints, for example, $a = 0 \wedge b \geq 1$ as $\{a = 0, b \geq 1\}$.

Definition 1


For an atom $p(\mathbf{x})$ and a set of CHCs P we write $P \vdash_{\mathbb{T}} p(\mathbf{x})$ if there exists a feasible AND-tree with root labelled by $p(\mathbf{x})$.

The soundness and completeness of derivation trees (Jaffar et al. 1998) implies that P is inconsistent if and only if $P \vdash_{\mathbb{T}} \mathbf{false}$.

On the right is an AND-tree corresponding to the derivations of \mathbf{false} using the clauses $c6$ followed by $c4$, $c2$ and $c1$ from the program in Figure 1 (right).

Definition 2 (Initial clauses and nodes)

Let P be a set of CHCs, with a distinguished predicate p^I in P which we call the *initial predicate*. The *constrained facts* $\{(p^I(\mathbf{x}) \leftarrow \theta) \mid (p^I(\mathbf{x}) \leftarrow \theta) \in P\}$ are called the *initial clauses* of P . Let t be an AND-tree for P . A node labelled by an identifier of the clause $p^I(\mathbf{x}) \leftarrow \theta$ is an *initial node* of t . We extend the term “initial predicate” and use the symbol p^I to refer also to renamed versions of the initial predicate that arise during clause transformations.



3 Precondition Inference

This section describes an approach to precondition generation. We limit our attention to sets of clauses for which every AND-tree for \mathbf{false} (whether feasible or infeasible) has at least one initial node. Although it is not decidable for an arbitrary set of CHCs P whether every derivation of \mathbf{false} uses the initial predicate, the above condition on AND-trees can be checked syntactically from the predicate dependency graph for P .

Definition 3 (Safe precondition)

Let P be a set of CHCs. Let ϕ be a constraint over \mathbb{T} , and let P' be the set of clauses obtained from P by replacing the initial clauses $\{(p^I(\mathbf{x}) \leftarrow \theta_i) \mid 1 \leq i \leq k\}$ by $\{(p^I(\mathbf{x}) \leftarrow \theta_i \wedge \phi) \mid 1 \leq i \leq k\}$. Then ϕ is a *safe precondition* for P if $P' \not\vdash_{\mathbb{T}} \mathbf{false}$.

Thus a safe precondition is a constraint that, when conjoined with the constraints on the initial predicate, is sufficient to block derivations of \mathbf{false} (given that we assume clauses for which p^I is essential for any derivation of \mathbf{false}).

Ideally we would like to find the most general, or *weakest* safe precondition. It is not computable in general, so we aim to find a condition that is as weak as possible. The constraint \mathbf{false} is always a safe precondition, albeit an uninteresting one. On the other hand, if $P \not\vdash_{\mathbb{T}} \mathbf{false}$ then any constraint, including \mathbf{true} , is a safe precondition for P .

We first show how a safe precondition can be derived from a set of clauses.

Definition 4 (Safe precondition $\mathbf{presafe}(P)$ extracted from a set P of clauses)

Let P be a set of clauses. The safe precondition $\mathbf{presafe}(P)$ is defined as:

$$\mathbf{presafe}(P) = \neg \bigvee \{\theta \mid (p^I(\mathbf{x}) \leftarrow \theta) \in P\}.$$

$\mathbf{presafe}(P)$ is clearly a safe precondition for P since for each initial clause $p^I(\mathbf{x}) \leftarrow \theta$ the conjunction $\mathbf{presafe}(P) \wedge \theta$ is \mathbf{false} . This precondition trivially blocks any derivation of \mathbf{false} since we assume that every derivation of \mathbf{false} uses an initial

clause. We next show how to construct a sequence P_0, P_1, \dots, P_m where $P = P_0$ and each element of the sequence is more specialised with respect to derivations of **false**, and as a consequence, the constraints in the initial clauses are stronger. Applying Definition 4 to P_m thus yields a weaker safe precondition for P .

3.1 Specialisation of Clauses

Definition 5 (Specialisation transformation)

Let P be a set of clauses, and let A be an atom. We write $P \Longrightarrow_A P'$ for a *specialisation transformation* of P with respect to A , yielding a set of clauses P' , such that the following holds.

- $P \vdash_{\mathbb{T}} A$ if and only if $P' \vdash_{\mathbb{T}} A$; and
- if $(p^I(\mathbf{x}) \leftarrow \theta) \in P$ then there exists an initial clause $(p^I(\mathbf{x}) \leftarrow \phi) \in P'$ such that $\models_{\mathbb{T}} \phi \rightarrow \theta$.

Note that a specialisation requires not only that derivations of A are preserved, but also that the initial clauses are preserved and possibly strengthened.

Lemma 1

Let $P \Longrightarrow_{\mathbf{false}} P'$ be a specialisation transformation with respect to **false**. Then $\models_{\mathbb{T}} \text{presafe}(P) \rightarrow \text{presafe}(P')$.

Proof

This follows immediately from Definitions 4 and 5. \square

We now present specific transformations for CHCs that satisfy Definition 5. Applying these transformations enables the derivation of more precise safe preconditions. These are adapted from established techniques from the literature on CLP and Horn clause verification and analysis.

3.1.1 Specialising CHCs by Partial Evaluation (PE)

Partial evaluation (Jones et al. 1993) is a transformation that specialises a program with respect to a given input. The “input” for partial evaluation of a set of CHCs P is a (set of) constrained atom(s) $A \leftarrow \theta$. The result of partial evaluation is a set of CHCs P' preserving the derivations of every instance of A that satisfies θ , that is, $P' \vdash A\phi$ if and only if $P \vdash A\phi$ whenever $\mathbb{T} \models \theta\phi$. The partial evaluation algorithm described here is an instantiation of the “basic algorithm” for partial evaluation of logic programs in Gallagher (1993).

The basic algorithm can be presented as the computation of the limit of the increasing sequence S_0, S_1, S_2, \dots , where S_0 is the set of input constrained atoms and for $i \geq 0$, $S_{i+1} = S_0 \cup \text{abstract}_{\Psi}(\text{unfold}_P(S_i))$. The “unfolding rule” unfold_P and the abstraction operation abstract_{Ψ} are parameters of the algorithm. For the algorithm used in this paper, the unfolding rule $\text{unfold}_P(S)$ takes a set of constrained facts S , and “partially evaluates” each element of S , using the following procedure. For each $(p(\mathbf{x}) \leftarrow \theta) \in S$, first construct the set of clauses $p(\mathbf{x}) \leftarrow \psi' \wedge B'$ where

$p(\mathbf{x}) \leftarrow \psi \wedge B$ is a clause in P , and $\psi' \wedge B'$ is obtained by unfolding $\psi \wedge \theta \wedge B$ by selecting atoms so long as they are deterministic (atoms defined by a single clause) and is not a call to an initial predicate or a recursive predicate, and ψ' is satisfiable in \mathbb{T} . Unfolding with this rule is guaranteed to terminate; $\text{unfold}_P(S)$ returns the set of constrained facts $q(\mathbf{y}) \leftarrow \psi'|_{\mathbf{y}}$ where $q(\mathbf{y})$ is an atom in B' .

The abstraction operation abstract_Ψ ensures that the sequence S_0, S_1, \dots has a finite limit. It performs property-based abstraction (Graf and Saïdi 1997) of a set of constrained facts with respect to a finite set of properties Ψ (also a finite set of constrained facts). Then $\text{abstract}_\Psi(S)$ is defined as follows.

$$\begin{aligned} \text{abstract}_\Psi(S) &= \{\text{rep}_\Psi(p(\mathbf{x}) \leftarrow \theta) \mid (p(\mathbf{x}) \leftarrow \theta) \in S\}, \text{ where} \\ \text{rep}_\Psi(p(\mathbf{x}) \leftarrow \theta) &= p(\mathbf{x}) \leftarrow \bigwedge \{\psi \mid (p(\mathbf{x}) \leftarrow \psi) \in \Psi, \mathbb{T} \wedge \theta \models \psi\} \end{aligned}$$

The effect of $\text{abstract}_\Psi(S)$ is to generalise each $q(\mathbf{y}) \leftarrow \theta \in S$ to $q(\mathbf{y}) \leftarrow \psi$, where ψ is the conjunction of properties in Ψ that are implied by θ . Thus only a finite number of “versions” of $q(\mathbf{y})$ can be generated, ensuring that the size of the sets S_i is finite (at most $2^{|\Psi|}$). The larger Ψ is, the more versions can be produced. More versions could cause overhead without necessarily giving more specialisation; for example, several essentially identical definitions of predicates could be produced. Thus it is important to choose Ψ taking into account both precision and efficiency.

In the implemented algorithm, Ψ consists of the following constrained facts, generated from each clause $p(\mathbf{x}) \leftarrow \phi, p_1(\mathbf{x}_1), \dots, p_n(\mathbf{x}_n) \in P$.

- For $1 \leq i \leq n$, $p_i(\mathbf{x}_i) \leftarrow \phi|_{\mathbf{x}_i}$ and for each $z \in \mathbf{x}_i$, $p_i(\mathbf{x}_i) \leftarrow \phi|_{\{z\}}$.
- $p(\mathbf{x}) \leftarrow \phi|_{\mathbf{x}}$ and for each $z \in \mathbf{x}$, $p(\mathbf{x}) \leftarrow \phi|_{\{z\}}$.

The first set of constrained facts distinguishes different call contexts, while the second set distinguishes answers. Constraints on individual variables are extracted. This choice of Ψ was found by experiment to be a good compromise between precision and efficiency, but further experiment and analysis is needed.

Each S_i in the sequence gives rise to a set of clauses $\text{renameunfold}_{\Psi, P}(S_i)$, which applies the unfolding rule to each element of S_i and renames the predicates in the resulting clauses according to the different versions produced by abstract_Ψ . The predicate **false** is not renamed. The result returned by partial evaluation is $\text{renameunfold}_{\Psi, P}(S_k)$, where S_k is the limit of the sequence.

Example 1

Consider the partial evaluation of the clauses in Figure 1. $S_0 = \{\mathbf{false} \leftarrow \mathbf{true}\}$ and Ψ consists of the following nine constrained facts extracted from the clauses as explained above:

$$\left\{ \begin{array}{l} \mathbf{if}(A, B) \leftarrow A \geq 0. \mathbf{if}(A, B) \leftarrow A \geq 1. \mathbf{init}(A, B) \leftarrow A \leq 100. \\ \mathbf{init}(A, B) \leftarrow A \geq 101. \mathbf{while}(A, B) \leftarrow A \geq 0. \mathbf{while}(A, B) \leftarrow A \geq 1. \\ \mathbf{while}(A, B) \leftarrow A \leq 0 \wedge B = 0. \mathbf{while}(A, B) \leftarrow A \leq 0. \mathbf{while}(A, B) \leftarrow B = 0. \end{array} \right\}$$

Partial evaluation of the clauses generates the clauses R_0, R_1, \dots and sets of constrained facts S_0, S_1, \dots as shown in Figure 2.

i	S_i	$R_i = \text{renameunfold}_{\Psi, P}(S_i)$
0	$S_0 = \{\text{false} \leftarrow \text{true}\}$	$R_0 = \{\text{false} \leftarrow A \leq 0, B = 0, \text{while}_7(A, B)\}$
1	$S_1 = S_0 \cup \{\text{while}(A, B) \leftarrow A \leq 0, B = 0.\}$	$R_1 = R_0 \cup \{\text{while}_7(A, B) \leftarrow A \leq 0, B = 0, \text{if}_6(A, B), \text{while}_7(A, B) \leftarrow A = 0, B = 0, C = 1, D = 2, \text{while}_5(C, D)\}$
2	$S_2 = S_1 \cup \{\text{while}(A, B) \leftarrow A \geq 1, \text{if}(A, B) \leftarrow \text{true}\}$	$R_2 = R_1 \cup \{\text{while}_5(A, B) \leftarrow A \geq 1, \text{if}_2(A, B), \text{while}_5(A, B) \leftarrow A \geq 1, C - A = 1, D - B = 2, \text{while}_5(C, D), \text{if}_6(A, B) \leftarrow A \geq 0, A + C = 100, \text{init}_4(C, B), \text{if}_6(A, B) \leftarrow A \geq 1, C - A = 100, \text{init}_3(C, B)\}$
3	$S_3 = S_2 \cup \{\text{if}(A, B) \leftarrow A \geq 1, \text{init}(A, B) \leftarrow A \geq 101, \text{init}(A, B) \leftarrow A \leq 100.\}$	$R_3 = R_2 \cup \{\text{if}_2(A, B) \leftarrow A \geq 1, A + C = 100, \text{init}_1(C, B), \text{if}_2(A, B) \leftarrow A \geq 1, C - A = 100, \text{init}_3(C, B), \text{init}_4(A, B) \leftarrow A \leq 100, \text{init}_3(A, B) \leftarrow A \geq 101.\}$
4	$S_4 = S_3 \cup \{\text{init}(A, B) \leftarrow A \leq 99.\}$	$R_4 = R_3 \cup \{\text{init}_1(A, B) \leftarrow A \leq 99.\}$
5	$S_5 = S_4$	$R_5 = R_4$

Fig. 2: Steps performed during the run of partial evaluation

Note that three versions of the `init` predicate are generated (from the new constrained facts generated in steps 3 and 4), each having different constraints. As we will see in the next section, this allows the extraction of more precise preconditions for safety of the clauses than could be obtained from the original clauses.

Lemma 2

Partial evaluation using the procedure described above is a specialisation transformation (Definition 5).

Proof

The algorithm satisfies the standard condition of partial evaluation that it preserves derivations of the given goal atom. The strengthening of the initial clauses follows from the fact that our unfolding rule does not unfold the initial predicate. Hence the result contains the initial clauses from the original, with constraints possibly strengthened by the call constraints in the algorithm. (If a clause is never called, its constraint is strengthened to *false*). \square

The safe precondition of the partially evaluated clauses is $\neg(A \leq 99 \vee A \leq 100 \vee A \geq 101)$, which is equivalent to *false* (over the integers). Thus partial evaluation has not improved the safe precondition compared to the original clauses in Figure 1. However, the splitting of the initial clauses enables a further specialisation, which is described next.

$\begin{aligned} &\text{false} \leftarrow A \geq 0, p(A, B). \\ &p(A, B) \leftarrow C \geq A, p(C, B). \\ &p(A, B) \leftarrow A = B. \end{aligned}$	$\begin{aligned} &\text{false} \leftarrow A \geq 0, \underline{B \geq A}, \underline{A \geq 0}, p(A, B). \\ &p(A, B) \leftarrow C \geq A, \underline{B \geq C}, \underline{C \geq 0}, p(C, B). \\ &p(A, B) \leftarrow A = B, \underline{B \geq A}, \underline{A \geq 0}. \end{aligned}$
---	---

Fig. 3: Example program (left) and its constraint specialised version (right)

3.1.2 Transforming CHCs by Constraint Specialisation (CS)

Constraint specialisation is a transformation that strengthens the constraints in a set of CHCs, while preserving derivations of a given atom. Consider the following simple example in Figure 3 (left) that motivates the principles of the transformation. Assume we wish to preserve derivations of **false**. The transformation in Figure 3 (right) is a constraint specialisation with respect to **false**. The strengthened constraints are obtained by recursively propagating $A \geq 0$ top-down from the goal **false** and $A = B$ bottom-up from the constrained fact. An invariant $B \geq A, A \geq 0$ for the derived answers of the recursive predicate $p(A, B)$ in derivations of **false** is computed and conjoined to each call to p in the clauses (underlined in the clauses in Figure 3 (right)).

Definition 6 (Constraint specialisation)

A constraint specialisation of P with respect to a goal A is a transformation in which each constraint ϕ in a clause of P is replaced by a constraint ψ where $\models_{\mathbb{T}} \psi \rightarrow \phi$, such that the resulting set of clauses is a specialisation transformation (Definition 5) of P with respect to A .

In our experiments, the combined top-down and bottom-up propagation of constraints illustrated above is achieved by abstract interpretation over the domain of convex polyhedra applied to a query-answer transformed version of the set of CHCs. The method is described in detail in Kafle and Gallagher (2017a). The result of applying constraint specialisation to the output of partial evaluation of the running example is shown in Figure 4. Note that the second clause for `if_6` has been eliminated, since its constraint was specialised to *false*.

The safe precondition derived after constraint specialisation from the initial clauses in Figure 4 is

$$\neg((A = 100 \wedge B = 0) \vee (A \leq 99 \wedge 2A + B = 200) \vee (A \geq 101 \wedge 2A - B = 200))$$

This simplifies (over the integers) to $B \neq |2A - 200|$, which is the condition obtained in Section 1 and is optimal (weakest).

3.1.3 Transforming CHCs by Trace Elimination (TE)

Let P be a set of CHCs and let t be an AND-tree for P . It is possible to construct a set of clauses P' which preserves the set of AND-trees (modulo predicate renaming) of P , apart from t . The transformation from P to P' is called *trace elimination* (of t). We have previously described a technique for trace elimination (Kafle and Gallagher 2017b), based on the difference operation on finite tree automata. In that work, trace elimination played the role of a refinement operation, in which infeasible

```

false ← A = 0, B = 0, while_7(A, B).  while_7(A, B) ← A = 0, B = 0, if_6(A, B).
while_7(A, B) ← A = 0, B = 0, C = 1, D = 2, while_5(C, D).
if_6(A, B) ← A = 0, B = 0, C = 100, init_4(C, B).
while_5(A, B) ← A ≥ 1, 2A - B = 0, if_2(A, B).
while_5(A, B) ← A ≥ 1, 2A = B, C - A = 1, D - 2A = 2, while_5(C, D).
if_2(A, B) ← A ≥ 1, 2A = B, A + C = 100, init_1(C, B).
if_2(A, B) ← A ≥ 1, 2A = B, C - A = 100, init_3(C, B).
init_4(A, B) ← A = 100, B = 0.
init_3(A, B) ← A ≥ 101, 2A - B = 200.
init_1(A, B) ← A ≤ 99, 2A + B = 200.

```

Fig. 4: Constraint specialisation of the partially evaluated clauses in Figure 2

traces were removed from a set of CHCs in a counterexample-guided verification algorithm in the CEGAR style (Clarke et al. 2003).

For the purpose of deriving safe preconditions of a set of clauses P , we apply trace elimination to eliminate both infeasible and feasible AND-trees. AND-trees for **false** are obtained naturally from transformations such as partial evaluation or constraint specialisation. First consider the elimination of an infeasible AND-tree.

Lemma 3

Let P' be the result of eliminating an *infeasible* AND-tree t for **false** from P . Then $P \Rightarrow_{\text{false}} P'$.

Proof

All derivations of **false** are preserved, and the transformation generates only predicate-renamed copies of the original clauses, hence the initial clauses are preserved. \square

So in this case we have $\models_{\mathbb{T}} \text{presafe}(P) \rightarrow \text{presafe}(P')$. However, the elimination of a feasible AND-tree t for **false** is not as straightforward. Nevertheless, we can still use this transformation to derive safe preconditions, by the following lemma.

Lemma 4

Let P' be the result of eliminating a *feasible* AND-tree t for **false** from P . Let $p^I(\mathbf{x})$ be the atom label of an initial node of t and let $\theta = \text{constr}(t)|_{\mathbf{x}}$. Then $\text{presafe}(P) = \text{presafe}(P') \wedge \neg\theta$.

Proof

$\neg\theta$ is a sufficient condition, when conjoined with the body of the clause labelling the initial node, to make t infeasible. All other derivations of **false** from P are preserved in P' . Hence the conjunction of $\neg\theta$ and $\text{presafe}(P')$ is a safe precondition for P . \square

The usefulness of trace elimination is twofold. Firstly, it can cause splitting of the initial predicates, resulting in disjunctive pre-conditions. Secondly, the elimination of a feasible trace acts as a decomposition of the problem.

3.2 Inferring Weaker Preconditions

We can combine the various transformations to derive weaker preconditions, as shown in the following two propositions.

Proposition 1

Let $P = P_0$ and let the sequence P_0, P_1, \dots, P_m be a sequence such that $P_i \Longrightarrow_{\text{false}} P_{i+1}$ ($0 \leq i < m$). Then $\models_{\mathbb{T}} \text{presafe}(P) \rightarrow \text{presafe}(P_m)$.

Proof

By induction on the length of the sequence, applying Lemma 1. \square

If we also eliminate feasible traces, then we have to keep track of the substitutions arising from the eliminated trees.

Proposition 2

Let $P = P_0$, $\psi_0 = \text{true}$ and let the sequence $(P_0, \psi_0), (P_1, \psi_1), \dots, (P_m, \psi_m)$ be a sequence of pairs where for ($0 \leq i < m$)

- either $P_i \Longrightarrow_{\text{false}} P_{i+1}$ and $\psi_i = \psi_{i+1}$, or
- P_{i+1} is obtained by eliminating a feasible trace t from P_i , and $\psi_{i+1} = \psi_i \wedge \neg\theta$, where $\neg\theta$ is the constraint extracted from t , as in Lemma 4.

Then $\models_{\mathbb{T}} \text{presafe}(P) \rightarrow (\text{presafe}(P_m) \wedge \psi_m)$.

Proof

By induction on the length of the sequence, applying Lemma 1 and Lemma 4. \square

Proposition 2 establishes the correctness of the algorithm used in Section 4, and any other algorithm that applies partial evaluation, constraint specialisation and trace elimination in any order. Proposition 1 is a special case of Proposition 2: if we do not eliminate any feasible trees then ψ_m is *true* and so $\models_{\mathbb{T}} \text{presafe}(P) \rightarrow \text{presafe}(P_m)$.

As we have shown, applying partial evaluation followed by constraint specialisation for our running example was sufficient to derive the weakest safe precondition. However, in more complex cases we need one or more iterations of these operations, possibly with the elimination of feasible AND-trees as well. In Figure 5 we show an example taken from Beyer et al. (2007) in which repeated application of partial evaluation followed by constraint specialisation does not achieve a useful result, but where the elimination of a single feasible AND-tree causes an optimal precondition to be generated. The optimal precondition for this program is $\text{init}(\mathbf{I}, \mathbf{A}, \mathbf{B}, \mathbf{N}) \leftarrow \mathbf{N} \leq \mathbf{I} \wedge \mathbf{A} + \mathbf{B} = 3 * \mathbf{N}$. To derive this, one needs to propagate constraints from the third and the fourth clauses (constrained facts corresponding to the predicate `1`) to the `init` clause. Since these constraints are disjunctive (arising from two different clauses), the propagation should be able to split the `init` predicate. PE can often perform splitting but not in this case since the recursive predicate `1` is not unfolded, owing to the potential for a resulting blowup.

We now show how trace-elimination together with other transformations allows us to derive this condition. Applying CS followed by PE to Figure 5 gives us the

```

false ← init(I, A, B, N), l(I, A, B, N).
l(I, A, B, N) ← I < N, l_body(A, B, A1, B1), I1 = I + 1, l(I1, A1, B1, N).
l(I, A, B, N) ← I ≥ N, A + B > 3 * N.
l(I, A, B, N) ← I ≥ N, A + B < 3 * N.
l_body(A0, B0, A1, B1) ← A1 = A0 + 1, B1 = B0 + 2.
l_body(A0, B0, A1, B1) ← A1 = A0 + 2, B1 = B0 + 1.
init(I, A, B, N).

```

Fig. 5: Example requiring trace elimination

```

c1. false ← init(A, B, C, D), l_3(A, B, C, D).
c2. l_3(A, B, C, D) ← -C + F >= 1, -A + D > 0, C - F >= -2, A - E = -1,
    B + C - F - G = -3, l_body_2(B, C, G, F), l_1(E, G, F, D).
c3. l_3(A, B, C, D) ← B + C - 3 * D > 0, A - D >= 0.
c4. l_3(A, B, C, D) ← -B - C + 3 * D > 0, A - D >= 0.
c5. l_1(A, B, C, D) ← -C + F >= 1, -A + D > 0, C - F >= -2, A - E = -1,
    B + C - F - G = -3, l_body_2(B, C, G, F), l_1(E, G, F, D).
c6. l_1(A, B, C, D) ← B + C - 3 * D > 0, -A + D > -1, A - D >= 0.
c7. l_1(A, B, C, D) ← -B - C + 3 * D > 0, -A + D > -1, A - D >= 0.
c8. l_body_2(A, B, C, D) ← A - C = -1, B - D = -2.
c9. l_body_2(A, B, C, D) ← A - C = -2, B - D = -1.
c10. init(A, B, C, D).

```

Fig. 6: The constraint specialisation of the program in Figure 5

program in Figure 6 (we have labelled the clauses for the purpose of presentation). If we derive a precondition from this program, we will get trivial *false*. As a next step, we search for a derivation (counterexample) violating safety. The trace tree $c1(c10, c2(c8, c5(c8, c5(c8, c5(c8, c6))))))$ (using its term representation) is a feasible counterexample. Then we remove this from the program in Figure 6 using the automata-theoretic approach described by Kafle and Gallagher (2017b). In summary, the approach consists of representing the program as well as the trace to be removed as finite tree automata, performing automata difference and generating a new program from the difference automaton. The new program is guaranteed not to contain the particular trace any more.

The removal causes the splitting of the predicate *l*, which the partial evaluation can take advantage of in the next iteration. Re-application of PE followed by CS generates the following clauses for *init* predicates (other clauses are not shown).

```

init_1(A, C, D, B) ← B > A.
init_2(A, C, D, B) ← A >= B, C + D > 3B.
init_3(A, C, D, B) ← A >= B, 3 * B > C + D.

```

Then the derived safe precondition is:

$$\text{init}(A, C, D, B) \leftarrow \neg((B > A) \vee (A \geq B \wedge C + D > 3B) \vee (A \geq B \wedge 3 * B > C + D)).$$

Simplifying the formula and mapping to the original variables, we obtain the following formula as the final precondition

$$\text{init}(I, A, B, N) \leftarrow N \leq I \wedge A + B = 3 * N.$$

There is, however, a performance-precision trade-off when removing (in)feasible

AND-trees. Trace elimination helps derive precise preconditions at the cost of performance; the *Fischer* protocol is an example of this. It requires 4 iterations of PE followed by CS to generate the optimal precondition (obtained in ≈ 8 seconds), whereas these iterations interleaved by trace elimination require only 3 iterations (but obtained in ≈ 35 seconds).

4 Experimental Evaluation

4.1 Benchmarks

We have experimented with three kinds of benchmarks.

1. *Unsafe I*: Examples that are known to be unsafe, where the initial states are over-general. In such cases the aim of safe precondition generation is to find out whether there is a useful subset of the initial states that is safe.
2. *Unsafe II*: Examples that are known to be unsafe, where the initial state is a counterexample state from which `false` can be derived. In this case it is pointless to try to find a safe subset as above, so we remove the given constraint on the initial state, and then try to derive a non-trivial safe precondition.
3. *Safe*: Examples that are safe for given initial states. In such cases, our aim is to try to weaken the conditions on the initial states. This is done by removing the given constraints from the initial states and then deriving safe preconditions. If we can generate safe preconditions that are more general than the original constraints then we have generalised the program without losing safety.

For the experiments, we collected a set of 241 (188 safe/53 unsafe) programs from a variety of sources. Most are from the repositories of state-of-the-art software verification tools such as DAGGER¹ (Gulavani et al. 2008), TRACER² (Jaffar et al. 2012), InvGen³ (Gupta and Rybalchenko 2009), and from the TACAS 2013 Software Verification Competition (Beyer 2013, Control flow and Loops categories).⁴ Other examples are from the literature on precondition generation, backwards analysis and parameter synthesis (Bakhirkin et al. 2014; Miné 2012a; Miné 2012b; Moy 2008; Bakhirkin and Monniaux 2017; Cassez et al. 2017) and manually translated to CHCs. These benchmarks are designed to demonstrate/test the strengths/usability of different tools and methods proposed to solve software verification, parameter synthesis and precondition generation problems and contain up to approximately 500 lines of code. Finally there are examples crafted by us; these are simple but non-trivial examples whose optimal precondition can be derived manually.

4.2 Implementation

We implemented an algorithm that builds a sequence as defined in Proposition 2, of length $3n + 2$ ($n \geq 0$), iteratively applying the transformations *pe* (partial evalua-

¹ <http://www.cfdvs.iitb.ac.in/~bhargav/dagger.php>

² <https://github.com/tracer-x/tracer/tree/master/test/transformation>

³ <http://www.mpi-sws.org/~agupta/invgen>

⁴ Translated to CHCs using the program specialisation approach of De Angelis et al. (2017).

	$n = 0$	$n = 1$	$n = 2$	$n = 3$
Safe instances (188)				
non-trivial (more general)	119 (101)	143 (125)	156 (129)	160 (131)
trivial/timeouts	69/0	45/3	32/10	28/16
avg. time (sec.)	1.45	14.69	27.52	36.73
Unsafe I instances (17)				
non-trivial	16	17	17	17
trivial/timeouts	1/0	0/0	0/0	0/0
avg. time (sec.)	0.23	0.82	1.64	3.35
Unsafe II instances (36)				
non-trivial	9	12	12	12
trivial/timeouts	27/0	24/2	24/7	24/7
avg. time (sec.)	3.38	50.41	64.72	70.91

Table 1: Results on 241 (188 safe and 53 unsafe) programs; timeout 5 minutes

tion), *cs* (constraint specialisation) and *te* (trace elimination). The safe precondition for P is $\text{presafe}(cs \circ pe \circ (te \circ cs \circ pe)^n(P))$ ($n \geq 0$). This particular sequence of transformations is based on the rationale that constraint specialisation is most effective when performed just after partial evaluation, which propagates constraints and introduces new versions of predicates. Trace elimination is more expensive and is performed only after the first iteration. In future work we will experiment with other strategies, especially to limit the application of *te*. The implementation is based on components from the RAHFT verifier (Kafle et al. 2016). This accepts CHCs (over the background theory of linear arithmetic) as input and returns a Boolean combination of linear constraints in terms of the initial state variables as a precondition. The tool is written in Ciao Prolog (Hermenegildo et al. 2012) and uses Yices 2.2 (Dutertre 2014) and the Parma Polyhedra Library (Bagnara et al. 2008) for constraint manipulation. The experiments were carried out on a MacBook Pro with a 2.7 GHz Intel Core i5 processor and 16 GB memory running OS X 10.11.6, with a timeout of 300 seconds for each example.

4.3 Discussion

Experimental results are shown in Table 1, for varying number of specialisation iterations n . The classifications “more general” and “non-trivial” in Table 1 relate the derived precondition I with the original condition on the initial states O . If $\models_{\top} I \not\equiv \text{false}$ then the result is non-trivial. If $\models_{\top} O \rightarrow I$ then the derived precondition is more general than the given initial states. For the *safe* benchmarks, the “more general” results are a subset of the “non-trivial” results, while for the *unsafe* benchmarks, the result cannot be more general than the original (unsafe) condition and so there are no “more general” results. *Timeouts* indicates the number of timeouts in the current iteration. When there is a timeout in the current iteration, the precondition is the precondition generated in the previous iteration. Therefore, the timeouts in the current iteration correspond to trivial, non-trivial or timeouts in the previous iteration. Thus, the trivial instances in the current iteration is the sum of trivial instances in this iteration and the trivial instances in the previous iteration of the current timeouts.

The choice of 3 iterations is motivated by the following observations (though we can stop at any iteration and still derive a precondition): (i) for the categories literature and hand-crafted benchmarks, 3 iterations suffice to reproduce earlier results, and (ii) iterations beyond the third yield negligible improvements but more timeouts.

For the *safe* benchmarks, the algorithm succeeds for $n = 3$ in generalising the safe initial conditions in 131 of the 188 benchmarks, and returns a non-trivial safe precondition in 160 of them. The remainder either return trivial results or a timeout. A higher proportion of the *unsafe* benchmarks return a trivial safe precondition, even when the initial state constraints are removed. A possible reason is that some of these unsafe programs are designed with an internal bug, and thus have no safe initial states. If the analysis returns a trivial safe precondition, it could be due to imprecision of the analysis, but could also be an indication to the programmer to look for the problem elsewhere than in the initial states.

The results in the column $n = 0$ show that the specialisation ($cs \circ pe$) alone can infer non-trivial preconditions for a large number of benchmarks, namely 63% (safe) and 37% (unsafe) instances both in less than 10 seconds. Among 119 non-trivial safe instances, 101 are generalised constraints.

Further specialisation ($n > 0$) increases the number of non-trivial and generalised preconditions by relatively small percentages of the total. The increased precision of the preconditions comes at a significant cost in time. For Safe, Unsafe I, and Unsafe II instances, the average time goes from 1.45, 0.23 and 3.38 seconds, respectively, when $n = 0$, to 36.73, 3.35 and 70.91 seconds, when $n = 3$. However, our prototype implementation is amenable to much optimisation, including sharing results from one iteration to the next, which could reduce the overhead.

For the categories of literature and hand-crafted benchmarks in which we know the weakest safe precondition, the tool is able to reproduce the results from the literature, see Figure 7. The results were generated in at most 1 iteration in less than a second, except for the *Fischer protocol*, which required 3 iterations and 35 seconds. As well as reproducing challenging examples from the literature (Figure 7), we are able to apply the technique to larger examples (shown in Table 1) than have previously been dealt with by automatic methods for precondition generation; we are also able to solve challenging examples that were not solvable by previous automatic techniques (such as our running example from Figure 1).

5 Concluding Remarks

We have presented a framework for computing a sufficient precondition of a program with respect to assertions; it enables derivation of more precise preconditions through iteration. Rather than relying on weakest precondition calculation or intricate transfer functions, it uses off-the-shelf components from program transformation and abstract interpretation, which eases implementation. Furthermore, the approach does not depend on specific abstract domain properties such as pseudo-complementation but is domain-independent and generic. By this we mean that the individual specialisation transformations such as partial evaluation and constraint

Program	Precondition
bakhirkin-fig3 (Bakhirkin et al. 2014)	$(1 \leq a \leq 99 \rightarrow b \geq 1) \wedge (a \leq 0 \rightarrow b \neq 0)$
bakhirkin (Bakhirkin et al. 2014)	$1 \leq a \leq 60 \vee a \geq 100$
mine (Miné 2012a)	$0 \leq a \leq 5$
mon_fig1 (Bakhirkin and Monniaux 2017)	$a = b \wedge a \geq 0$
moy (Moy 2008)	$b < 1 \vee (b < 2 \wedge a > 0)$
navas2 (crafted)	$a \leq 99 \vee b \geq 100$
simple_function (Miné 2012b)	$6 \leq a \leq 61$
test_both_branches (Miné 2012b)	$3 \leq a \leq 17$
test_nondet_body (Miné 2012b)	$6 \leq a \leq 13$
test_nondet_cond (Miné 2012b)	$3 \leq a \leq 17$
test_then_branch (Miné 2012b)	$10 \leq a \leq 20$
fischer (Cassez et al. 2017)	$a + 2c < b \vee a < 0 \vee b < 0 \vee c \leq 0$
Jhala (Jhala and McMillan 2006)	$a < 0 \vee a \geq b \vee c \neq d$
Ball SLAM (Ball et al. 2004)	$b < c$
client ssh protocol	$b < a \vee b < 2 \vee a > 3$
Beyer et al. (2007)	$n \leq i \wedge a + b = 3n$

Fig. 7: Examples and their safe preconditions

specialisation can be adapted to different abstract domains with their usual precision/performance limits, while still using features of the framework such as iteration and disjunctive constraints that arise from polyvariant specialisation. Evaluation on a set of benchmarks is promising. We are currently investigating the conditions under which the derived preconditions are the weakest possible, as well as other improved termination criteria for refinement with the aim of generating optimal preconditions.

Acknowledgements

We are grateful for support from the Australian Research Council. The work has been supported by Discovery Project grant DP140102194, and Graeme Gange is supported through Discovery Early Career Researcher Award DE160100568. We wish to thank Jorge Navas, for useful discussions based on an early draft of the manuscript, Emanuele De Angelis, for making benchmarks available to us, and the three anonymous reviewers, for suggestions which led to clear improvements of the paper.

References

- BAGNARA, R., HILL, P. M., AND ZAFFANELLA, E. 2008. The Parma Polyhedra Library: Toward a complete set of numerical abstractions for the analysis and verification of hardware and software systems. *Sci. Comput. Program.* 72, 1–2, 3–21.
- BAKHIRKIN, A., BERDINE, J., AND PITERMAN, N. 2014. Backward analysis via over-approximate abstraction and under-approximate subtraction. In *Static Analysis*, M. Müller-Olm and H. Seidl, Eds. LNCS, vol. 8723. Springer, 34–50.
- BAKHIRKIN, A. AND MONNIAUX, D. 2017. Combining forward and backward abstract interpretation of Horn clauses. In *Static Analysis*, F. Ranzato, Ed. LNCS, vol. 10422. Springer, 23–45.
- BALL, T., COOK, B., LEVIN, V., AND RAJAMANI, S. K. 2004. SLAM and static driver

- verifier: Technology transfer of formal methods inside Microsoft. In *Integrated Formal Methods*, E. A. Boiten, J. Derrick, and G. Smith, Eds. LNCS, vol. 2999. Springer, 1–20.
- BEYER, D. 2013. SV-COMP 2013. In *Tools and Algorithms for the Construction and Analysis of Systems*, N. Piterman and S. A. Smolka, Eds. LNCS, vol. 7795. Springer, 594–609.
- BEYER, D., HENZINGER, T. A., MAJUMDAR, R., AND RYBALCHENKO, A. 2007. Path invariants. In *Programming Language Design and Implementation*, J. Ferrante and K. S. McKinley, Eds. ACM, 300–309.
- CASSEZ, F., JENSEN, P. G., AND LARSEN, K. G. 2017. Refinement of trace abstraction for real-time programs. In *Reachability Problems*, M. Hague and I. Potapov, Eds. LNCS, vol. 10506. Springer, 42–58.
- CLARKE, E. M., GRUMBERG, O., JHA, S., LU, Y., AND VEITH, H. 2003. Counterexample-guided abstraction refinement for symbolic model checking. *J. ACM* 50, 5, 752–794.
- COUSOT, P. AND COUSOT, R. 1992. Abstract interpretation and application to logic programs. *J. Log. Program.* 13, 2&3, 103–179.
- COUSOT, P., COUSOT, R., AND LOGOZZO, F. 2011. Precondition inference from intermittent assertions and applications to contracts on collections. In *Verification, Model Checking and Abstract Interpretation*, R. Jhala and D. Schmidt, Eds. LNCS, vol. 6538. Springer, 150–168.
- DE ANGELIS, E., FIORAVANTI, F., PETTOROSSO, A., AND PROIETTI, M. 2014. Program verification via iterated specialization. *Sci. Comput. Program.* 95, 149–175.
- DE ANGELIS, E., FIORAVANTI, F., PETTOROSSO, A., AND PROIETTI, M. 2017. Semantics-based generation of verification conditions via program specialization. *Sci. Comput. Program.* 147, 78–108.
- DUTERTRE, B. 2014. Yices 2.2. In *Computer-Aided Verification*, A. Biere and R. Bloem, Eds. LNCS, vol. 8559. Springer, 737–744.
- GALLAGHER, J. P. 1993. Specialisation of logic programs: A tutorial. In *Proc. ACM SIGPLAN Symp. PEPM'93*. ACM Press, 88–98.
- GALLAGHER, J. P. AND LAFAVE, L. 1996. Regular approximation of computation paths in logic and functional languages. In *Partial Evaluation*, O. Danvy, R. Glück, and P. Thiemann, Eds. LNCS, vol. 1110. Springer, 115–136.
- GRAF, S. AND SAÏDI, H. 1997. Construction of abstract state graphs with PVS. In *Computer Aided Verification, 9th International Conference, CAV*, O. Grumberg, Ed. Lecture Notes in Computer Science, vol. 1254. Springer, 72–83.
- GREBENSHCHIKOV, S., LOPES, N. P., POPEEA, C., AND RYBALCHENKO, A. 2012. Synthesizing software verifiers from proof rules. In *Proc. 33rd ACM SIGPLAN Conf. Programming Language Design and Implementation*, J. Vitek, H. Lin, and F. Tip, Eds. ACM, 405–416.
- GULAVANI, B. S., CHAKRABORTY, S., NORI, A. V., AND RAJAMANI, S. K. 2008. Automatically refining abstract interpretations. In *Tools and Algorithms for the Construction and Analysis of Systems*, C. R. Ramakrishnan and J. Rehof, Eds. LNCS, vol. 4963. Springer, 443–458.
- GUPTA, A. AND RYBALCHENKO, A. 2009. InvGen: An efficient invariant generator. In *Computer-Aided Verification*, A. Bouajjani and O. Maler, Eds. LNCS, vol. 5643. Springer, 634–640.
- GURFINKEL, A., KAHSAI, T., KOMURAVELLI, A., AND NAVAS, J. A. 2015. The Seahorn verification framework. In *Computer Aided Verification*, D. Kroening and C. S. Păsăreanu, Eds. LNCS, vol. 9206. Springer, 343–361.
- HERMENEGILDO, M. V., BUENO, F., CARRO, M., LÓPEZ-GARCÍA, P., MERA, E., MORALES, J. F., AND PUEBLA, G. 2012. An overview of ciao and its design philosophy. *TPLP* 12, 1-2, 219–252.

- HOWE, J. M., KING, A., AND LU, L. 2004. Analysing logic programs by reasoning backwards. In *Program Development in Computational Logic*, M. Bruynooghe and K. Lau, Eds. LNCS, vol. 3049. Springer, 152–188.
- JAFFAR, J., MAHER, M., MARRIOTT, K., AND STUCKEY, P. J. 1998. The semantics of constraint logic programs. *J. Log. Program.* 37, 1–3, 1–46.
- JAFFAR, J., MURALI, V., NAVAS, J. A., AND SANTOSA, A. E. 2012. TRACER: A symbolic execution tool for verification. In *Computer-Aided Verification*, P. Madhusudan and S. A. Seshia, Eds. LNCS, vol. 7358. Springer, 758–766.
- JHALA, R. AND MCMILLAN, K. L. 2006. A practical and complete approach to predicate refinement. In *Tools and Algorithms for the Construction and Analysis of Systems*, H. Hermanns and J. Palsberg, Eds. LNCS, vol. 3920. Springer, 459–473.
- JONES, N., GOMARD, C., AND SESTOFT, P. 1993. *Partial Evaluation and Automatic Software Generation*. Prentice Hall.
- KAFLE, B. AND GALLAGHER, J. P. 2017a. Constraint specialisation in Horn clause verification. *Sci. Comput. Program.* 137, 125–140.
- KAFLE, B. AND GALLAGHER, J. P. 2017b. Horn clause verification with convex polyhedral abstraction and tree automata-based refinement. *Computer Languages, Systems & Structures* 47, 2–18.
- KAFLE, B., GALLAGHER, J. P., AND MORALES, J. F. 2016. RAHFT: A tool for verifying Horn clauses using abstract interpretation and finite tree automata. In *Computer-Aided Verification*, S. Chaudhuri and A. Farzan, Eds. LNCS, vol. 9779. Springer, 261–268.
- MARRIOTT, K. AND SØNDERGAARD, H. 1993. Precise and efficient groundness analysis for logic programs. *ACM LOPLAS* 2, 1–4, 181–196.
- MINÉ, A. 2006. The octagon abstract domain. *Higher-Order and Symbolic Computation* 19, 1, 31–100.
- MINÉ, A. 2012a. Inferring sufficient conditions with backward polyhedral under-approximations. *Electr. Notes Theor. Comput. Sci.* 287, 89–100.
- MINÉ, A. 2012b. Sufficient condition polyhedral prototype analyzer. <https://www-apr.lip6.fr/~mine/banal/syntax.html>. Accessed: 2018-01-23.
- MOY, Y. 2008. Sufficient preconditions for modular assertion checking. In *Verification, Model Checking and Abstract Interpretation*, F. Logozzo, D. A. Peled, and L. D. Zuck, Eds. LNCS, vol. 4905. Springer, 188–202.
- PERALTA, J. C., GALLAGHER, J. P., AND SAĞLAM, H. 1998. Analysis of imperative programs through analysis of constraint logic programs. In *Static Analysis*, G. Levi, Ed. LNCS, vol. 1503. 246–261.