Convex polyhedral abstractions, specialisation and property-based predicate splitting in Horn clause verification

Kafle, Bishoksan; Gallagher, John Patrick

Convex polyhedral abstractions, specialisation and property-based predicate splitting in Horn clause verification*

Bishoksan Kafle
Roskilde University
Denmark
kafle@ruc.dk

John P. Gallagher
Roskilde University
Denmark
IMDEA Software Institute
Madrid, Spain
jpg@ruc.dk

We present an approach to constrained Horn clause (CHC) verification combining three techniques: abstract interpretation over a domain of convex polyhedra, specialisation of the constraints in CHCs using abstract interpretation of query-answer transformed clauses, and refinement by splitting predicates. The purpose of the work is to investigate how analysis and transformation tools developed for constraint logic programs (CLP) can be applied to the Horn clause verification problem. Abstract interpretation over convex polyhedra is capable of deriving sophisticated invariants and when used in conjunction with specialisation for propagating constraints it can frequently solve challenging verification problems. This is a contribution in itself, but refinement is needed when it fails, and the question of how to refine convex polyhedral analyses has not been studied much. We present a refinement technique based on interpolants derived from a counterexample trace; these are used to drive a property-based specialisation that splits predicates, leading in turn to more precise convex polyhedral analyses. The process of specialisation, analysis and splitting can be repeated, in a manner similar to the CEGAR and iterative specialisation approaches.

1 Introduction

In this paper we explore the use of techniques used in constraint logic program (CLP) analysis and specialisation, for the purpose of CHC verification. Pure CLP is syntactically and semantically the same as CHC. Unlike CLP, CHCs are not always regarded as executable programs, but rather as specifications or semantic representations of other formalisms. However these are only pragmatic distinctions and the semantic equivalence of CHC and CLP means that techniques developed in one framework are applicable to the other.

Relevant concepts from CLP include the approximation of the minimal model of a CLP program using abstract interpretation, specialisation of a CLP program with respect to a goal and model-preserving transformation of CLP programs. Relevant concepts drawn from the CHC verification literature include finding a model of a set of CHCs, property-based abstraction, counterexample generation, and refinement of property-based abstraction using interpolants.

The results shown in the paper are preliminary and much research remains to be done in exploiting the many connections and possibilities for cross-fertilisation between CLP and CHC. The contributions of this paper are:

- to demonstrate that abstract interpretation over convex polyhedra is capable of deriving sophisticated invariants, and when used in conjunction with specialisation for propagating constraints it can frequently solve challenging verification problems;

- to investigate the problem of refinement of polyhedral abstractions, drawing ideas from counterexample-guided refinement.

In Section 2 we define the basic notation and concepts needed for the verification procedure. Section 3 reviews the technique of abstract interpretation over convex polyhedra, applied to CLP/CHC, along with the important enhancement of this technique using widening thresholds. In Section 4 a procedure for specialisation of CHCs is described, based on query-answer transformations and abstract interpretation. A simple but surprisingly effective verification tool-chain combining specialisation with abstract interpretation is introduced. Section 5 explains how to use a (spurious) counterexample from a failed verification attempt to construct a property-based specialisation using interpolants. Experimental results and related works are reported in Section 6 and Section 7 respectively. Finally in Section 8 we conclude and discuss possible extensions and improvements.

2   Preliminaries

A CHC is a first order predicate logic formula of the form $\forall(\phi \wedge B_1(X_1) \wedge \ldots \wedge B_k(X_k) \rightarrow H(X))$ $(k \geq 0)$, where ϕ is a conjunction of constraints with respect to some background theory, X_i, X are (possibly empty) vectors of distinct variables, $B_1, \ldots, B_k, H$ are predicate symbols, $H(X)$ is the head of the clause and $\phi \wedge B_1(X_1) \wedge \ldots \wedge B_k(X_k)$ is the body. Sometimes the clause is written $H(X) \leftarrow \phi \wedge B_1(X_1), \ldots, B_k(X_k)$ and in concrete examples it is written in the form `H :- ` ϕ`, `B_1`(X`$_1$`),...,`B_k`(X`$_k$`)`. In examples, predicate symbols start with lowercase letters while we use uppercase letters for variables.

In this paper we take the constraint theory to be linear arithmetic with the relation symbols $\leq, \geq, <, >$ and $=$. There is a distinguished predicate symbol false which is interpreted as false. In practice the predicate false only occurs in the head of clauses; we call clauses whose head is false *integrity constraints*, following the terminology of deductive databases. Thus the formula $\phi_1 \leftarrow \phi_2 \wedge B_1(X_1), \ldots, B_k(X_k)$ is equivalent to the formula false $\leftarrow \neg\phi_1 \wedge \phi_2 \wedge B_1(X_1), \ldots, B_k(X_k)$. The latter might not be a CHC (e.g. if ϕ_1 contains $=$) but can be converted to an equivalent set of CHCs by transforming the formula $\neg\phi_1$ and distributing any disjunctions that arise over the rest of the body. For example, the formula `X=Y :- p(X,Y)` is equivalent to the set of CHCs `false :- X>Y, p(X,Y)` and `false :- X<Y, p(X,Y)`. Integrity constraints can be seen as safety properties. For example if a set of CHCs encodes the behaviour of a transition system, the bodies of integrity constraints represent unsafe states. Thus proving safety consists of showing that the bodies of integrity constraints are false in all models of the CHC clauses. Figure 1 shows an example set of CHCs (taken from [6]), modeled over reals containing an integrity constraint, and in this example the problem is to prove that the body of the first clause is unsatisfiable.

2.1   The CHC verification problem.

To state this more formally, given a set of CHCs P, the CHC verification problem is to check whether there exists a model of P. Obviously any model of P assigns false to the bodies of integrity constraints. We restate this property in terms of the derivability of the predicate false. Let $P \models F$ mean that F is a logical consequence of P, that is, that every interpretation satisfying P also satisfies F.

Lemma 1. *P has a model if and only if* $P \not\models$ false.

```
c1. false:- N>0,I=0,A=0,B=0, l(I,A,B,N).
c2. l(I,A,B,N):-I < N, l_body(A,B,A1,B1), I1 = I+1, l(I1,A1,B1,N).
c3. l(I,A,B,N):- I >=N, A + B > 3 * N.
c4. l(I,A,B,N):- I >=N, A + B < 3 * N.
c5. l_body(A0,B0,A1,B1):- A1 = A0+1, B1 = B0+2.
c6. l_body(A0,B0,A1,B1):- A1 = A0+2, B1 = B0+1.
```

Figure 1: Example program *t4.pl* [6]

Proof. Writing $I(F)$ to mean that interpretation I satisfies F, we have:

$$
\begin{aligned}
P \not\models \text{false} \quad &\equiv \quad \text{there exists some interpretation } I \text{ such that } I(P) \text{ and } \neg I(\text{false}) \\
&\qquad \textit{by definition of the } \models \textit{ relation} \\
&\equiv \quad \text{there exists some interpretation } I \text{ such that } I(P) \\
&\qquad (\textit{since } \neg I(\text{false}) \textit{ is true by defn. of } \text{false}) \\
&\equiv \quad P \text{ has a model.}
\end{aligned}
$$

$\square$

This lemma holds for arbitrary interpretations (only assuming that the predicate false is interpreted as false), uses only the textbook definitions of "interpretation" and "model" and does not depend on the constraint theory.

The verification problem can be formulated deductively rather than model-theoretically. We can exploit proof procedures for constraint logic programming [24] to reason about the satisfiability of a set of CHCs. Let the relation $P \vdash A$ denote that A is derivable from P using some proof procedure. If the proof procedure is sound then $P \vdash A$ implies $P \models A$, which means that $P \vdash \text{false}$ is a sufficient condition for P to have no model, by Lemma 1. This corresponds to using a sound proof procedure to find or check a counterexample. On the other hand to show that P does have a model, soundness is not enough since we need to establish $P \not\models \text{false}$. As we will see in Section 2.3 we approach this problem by using *approximations* to reason about the non-provability of false, applying the theory of abstract interpretation [11] to a complete proof procedure for atomic formulas (the "fixed-point semantics" for constraint logic programs [24, Section 4]). In effect, we construct by abstract interpretation a proof procedure that is *complete* (but possibly not sound) for proofs of atomic formulas. With such a procedure, $P \not\vdash \text{false}$ implies $P \not\models \text{false}$ and thus establishes that P has a model.

2.2 Representation of Interpretations

An interpretation of a set of CHCs is represented as a set of *constrained facts* of the form $A \leftarrow \mathscr{C}$ where A is an atomic formula $p(Z_1, \ldots, Z_n)$ where $Z_1, \ldots, Z_n$ are distinct variables and $\mathscr{C}$ is a constraint over $Z_1, \ldots, Z_n$. If $\mathscr{C}$ is true we write $A \leftarrow$ or just A. The constrained fact $A \leftarrow \mathscr{C}$ is shorthand for the set of variable-free facts $A\theta$ such that $\mathscr{C}\theta$ holds in the constraint theory, and an interpretation M denotes the set of all facts denoted by its elements; M assigns true to exactly those facts. $M_1 \subseteq M_2$ if the set of denoted facts of M_1 is contained in the set of denoted facts of M_2.

Minimal models. A model of a set of CHCs is an interpretation that satisfies each clause. There exists a minimal model with respect to the subset ordering, denoted $M[\![P]\!]$ where P is the set of CHCs. $M[\![P]\!]$

can be computed as the least fixed point (lfp) of an immediate consequences operator (called S_P^D in [24, Section 4]), which is an extension of the standard T_P operator from logic programming, extended to handle the constraint domain D. Furthermore $\mathsf{lfp}(S_P^D)$ can be computed as the limit of the ascending sequence of interpretations $\emptyset, S_P^D(\emptyset), S_P^D(S_P^D(\emptyset)), \ldots$. This sequence provides a basis for abstract interpretation of CHC clauses.

2.3   Proof Techniques

Proof by over-approximation of the minimal model.   It is a standard theorem of CLP that the minimal model $M[\![P]\!]$ is equivalent to the set of atomic consequences of P. That is, $P \models p(v_1, \ldots, v_n)$ if and only if $p(v_1, \ldots, v_n) \in M[\![P]\!]$. Therefore, the CHC verification problem for P is equivalent to checking that $\mathsf{false} \notin M[\![P]\!]$. It is sufficient to find a set of constrained facts M' such that $M[\![P]\!] \subseteq M'$, where $\mathsf{false} \notin M'$. This technique is called proof by *over-approximation of the minimal model*.

Proof by specialisation.   A specialisation of a set of CHCs P with respect to an atom A is the transformation of P to another set of CHCs P' such that $P \models A$ if and only if $P' \models A$. Specialisation is usually viewed as a program optimisation method, specialising some general-purpose program to a subset of its possible inputs, thereby removing redundancy and pre-computing statically determined computations. In our context we use specialisation to focus the verification problem on the formula to be proved. More specifically, we specialise a set of CHCs with respect to a "query" to the atom false; thus the specialised CHCs entail false if and only if the original clauses entailed false.

3   Abstract Interpretation over Convex Polyhedra

Convex polyhedron analysis (CPA) [12] is a program analysis technique based on abstract interpretation [11]. When applied to a set of CHCs P it constructs an over-approximation M' of the minimal model of P, where M' contains at most one constrained fact $p(X) \leftarrow \mathscr{C}$ for each predicate p. The constraint $\mathscr{C}$ is a conjunction of linear inequalities, representing a convex polyhedron. The first application of convex polyhedron analysis to CLP was by Benoy and King [4]. Since the domain of convex polyhedra contains infinite increasing chains, the use of a *widening* operator for convex polyhedra [11, 12] is needed to ensure convergence of the abstract interpretation. Furthermore much research has been done on improving the precision of widening operators. One techniques is known as widening-upto, or widening with thresholds [23]. A threshold is an assertion that is combined with a widening operator to improve its precision.

Recently, a technique for deriving more effective thresholds was developed [27], which we have adapted and found to be effective in experimental studies. The thresholds are computed by the following method. Let S_P^D be the standard immediate consequence operator for CHCs mentioned in Section 2.2. That is, if I is a set of constrained facts, $S_P^D(I)$ is the set of constrained facts that can be derived in one step from I. Given a constrained fact $p(\bar{Z}) \leftarrow \mathscr{C}$, define $\mathsf{atomconstraints}(p(\bar{Z}) \leftarrow \mathscr{C})$ to be the set of constrained facts $\{p(\bar{Z}) \leftarrow C_i \mid \mathscr{C} = C_1 \wedge \ldots \wedge C_k, 1 \leq i \leq k)\}$. The function $\mathsf{atomconstraints}$ is extended to interpretations by $\mathsf{atomconstraints}(I) = \bigcup_{p(\bar{Z}) \leftarrow \mathscr{C} \in I} \{\mathsf{atomconstraints}(p(\bar{Z}) \leftarrow \mathscr{C})\}$.

Let $I_\top$ be the interpretation consisting of the set of constrained facts $p(\bar{Z}) \leftarrow \mathsf{true}$ for each predicate p. We perform three iterations of S_P^D starting with $I_\top$ (the first three elements of a "top-down" Kleene sequence) and then extract the atomic constraints. That is, thresholds is defined as follows.

$$\mathsf{thresholds}(P) = \mathsf{atomconstraints}(S_P^{D(3)}(I_\top))$$

A difference from the method in [27] is that we use the concrete semantic function S_P^D rather than the abstract semantic function when computing thresholds. The set of threshold constraints represents an attempt to find useful predicate properties and when widening they help to preserve invariants that might otherwise be lost during widening. See [27] for further details. Threshold constraints that are not invariants are simply discarded during widening.

4 Specialisation by constraint propagation

We next present a procedure for specialising CHC clauses. In contrast to classical specialisation techniques based on partial evaluation with respect to a goal, the specialisation does not unfold the clauses at all; rather, we compute a specialised version of each clause in the program, in which the constraints from the goal are propagated top-down and answers are propagated bottom-up. The implementation is based on query-answer transformations and abstract interpretation over convex polyhedra.

Let P be a set of CHCs and let A be an atomic formula. For each clause $H \leftarrow \mathscr{B}$ in P we compute a new clause $H \leftarrow C, \mathscr{B}$ where C is a constraint, yielding a program P_A *specialised* for A. If the addition of C makes the clause body unsatisfiable, it is the same as removing the clause from P_A. Clearly P_A may have fewer consequences than P but our procedure guarantees that it preserves the inferability of (constrained instances of) A. That is, for every constraint C over the variables of A, $P \models \forall (C \rightarrow A)$ if and only if $P_A \models \forall (C \rightarrow A)$.

The procedure is as follows: the inputs are a set of CHCs P and an atomic formula A.

1. Compute a *query-answer transformation* of P with respect to A, denoted P_A^{qa}, containing predicates p^{q} and p^{a} for each predicate p in P.

2. Compute an over-approximation of the model of P_A^{qa}, expressed as a set of constrained facts $p^*(X) \leftarrow C$, where $*$ is q or a. We assume that each predicate p^* has exactly one constrained fact in the model (where C is possibly *false* or a disjunction).

3. For each clause $p(X) \leftarrow \mathscr{B}$ in P, let the model of p^{a} be $p^{\mathsf{a}}(X) \leftarrow C^{\mathsf{a}}$ (where X is the same tuple of variables in $p(X)$ and $p^{\mathsf{a}}(X)$).

4. Replace the clause $p(X) \leftarrow \mathscr{B}$ in P by $p(X) \leftarrow C^{\mathsf{a}}, \mathscr{B}$ in P_A.

Note that if for some predicate p, C^{a} is false, then all the clauses for p are removed in P_A as their bodies are unsatisfiable. We now explain each step in turn.

4.1 The query-answer transformation

The query-answer transformation was inspired by – but is a generalisation of – the magic-set transformation from deductive databases [3]. Its purpose, both in deductive databases and in subsequent applications in logic program analysis [15] was to simulate goal-directed (*top-down*) computation or deduction in a goal-independent (*bottom-up*) framework. Let us define the transformation.

Given a set of CHCs P and an atom A, the query-answer program for P wrt. A, denoted P_A^{qa}, consists of the following clauses. For an atom $A = p(t)$, A^{a} and A^{q} represent the atoms $p^{\mathsf{a}}(t)$ and $p^{\mathsf{q}}(t)$ respectively.

- (Answer clauses). For each clause $H \leftarrow C, B_1, \ldots, B_n$ $(n \geq 0)$ in P, P_A^{qa} contains the clause $H^{\mathsf{a}} \leftarrow C, H^{\mathsf{q}}, B_1^{\mathsf{a}}, \ldots, B_n^{\mathsf{a}}$.

- (Query clauses). For each clause $H \leftarrow C, B_1, \ldots, B_i, \ldots, B_n$ $(n \geq 0)$ in P, P_A^{qa} contains the following clauses:

$$B_1^{\mathsf{q}} \leftarrow C, H^{\mathsf{q}}.$$
$$\ldots$$
$$B_i^{\mathsf{q}} \leftarrow C, H^{\mathsf{q}}, B_1^{\mathsf{a}}, \ldots, B_{i-1}^{\mathsf{a}}.$$
$$\ldots$$
$$B_n^{\mathsf{q}} \leftarrow C, H^{\mathsf{q}}, B_1^{\mathsf{a}}, \ldots, B_{n-1}^{\mathsf{a}}.$$

- (Goal clause). $A^{\mathsf{q}} \leftarrow \mathsf{true}.$

The program P_A^{qa} encodes a left-to-right, depth-first computation of the query $\leftarrow A$ for CHC clauses P (that is, the standard CLP computation rule, SLD extended with constraints). This is a complete proof procedure, assuming that all clauses matching a given call are explored in parallel. (Note: the incompleteness of standard Prolog CLP proof procedures arises due to the fact that clauses are tried in a fixed order).

The relationship of the model of the program P_A^{qa} to the computation of the goal $\leftarrow A$ in P is expressed by the following property[1]. An SLD-derivation in CLP is a sequence $G_0, G_1, \ldots, G_k$ where each G_i is a goal $\leftarrow C, B_1, \ldots, B_m$, where C is a constraint and $B_1, \ldots, B_m$ are atoms. In a left-to-right computation, G_{i+1} is obtained by resolving B_1 with a program clause.

Property 1 (Correctness of query-answer transformation). *Let P be a set of CHCs and A be an atom. Let P_A^{qa} be the query-answer program for P wrt. A. Then*

(i) *if there is an SLD-derivation $G_0, \ldots, G_i$ where $G_0 = \leftarrow A$ and $G_i = \leftarrow C, B_1, \ldots, B_m$, then $P_A^{\mathsf{qa}} \models \forall (C|_{\mathsf{vars}(B_1)} \rightarrow B_1^{\mathsf{q}});$*

(ii) *if there is an SLD-derivation $G_0, \ldots, G_i$ where $G_0 = \leftarrow A$, containing a sub-derivation $G_{j_1}, \ldots, G_{j_k}$, where $G_{j_i} \leftarrow C', B_1, B'$ and $G_{j_k} = \leftarrow C, B'$, then $P_A^{\mathsf{qa}} \models \forall (C|_{\mathsf{vars}(B_1)} \rightarrow B_1^{\mathsf{a}})$. (This means that the atom B_1 in G_{j_i} was successfully answered, with answer constraint $C|_{\mathsf{vars}(B_1)}$).*

(iii) *As a special case of (ii), if there is a successful derivation of the goal $\leftarrow A$ with answer constraint C then $P_A^{\mathsf{qa}} \models \forall (C \rightarrow A^{\mathsf{a}})$.*

4.2   Over-approximation of the model of the query-answer program $P_{\mathsf{false}}^{\mathsf{qa}}$

The query-answer transformation of P with respect to false is computed. It follows from Property 1(iii) that if false is derivable from P then $P_{\mathsf{false}}^{\mathsf{qa}} \models \mathsf{false}^{\mathsf{a}}$. Convex polyhedral analysis of $P_{\mathsf{false}}^{\mathsf{qa}}$ yields an overapproximation of $M[\![P_{\mathsf{false}}^{\mathsf{qa}}]\!]$, say M', containing constrained facts for the query and answer predicates. These represent the calls and answers generated during all derivations starting from the goal false.

4.3   Strengthening the constraints in P

We use the information in M' to specialise the original clauses in P. Suppose M' contains constrained facts $p^{\mathsf{q}}(X) \leftarrow C^{\mathsf{q}}$ and $p^{\mathsf{a}}(X) \leftarrow C^{\mathsf{a}}$. If there is no constrained fact $p^*(X) \leftarrow C^*$ for some p^* then we consider M' to contain $p^*(X) \leftarrow false$. The clauses in P with head predicate p can be *strengthened* using the constraints C^{q} and C^{a}. Namely, for every clause $p(X) \leftarrow \mathscr{B}$ in P (assuming that the constrained facts are renamed to have the same variables X) the conjunction $C^{\mathsf{q}} \wedge C^{\mathsf{a}}$ are added to the body $\mathscr{B}$. The addition of C^{q} corresponds to propagating constraints "top-down" (via the calls) while the addition of

---

[1] Note that the model of P_A^{qa} might not correspond exactly to the calls and answers in the SLD-computation, since the CLP computation treats constraints as syntactic entities through decision procedures and the actual constraints could differ.

```
c1. false:- N>0,I=0,A=0,B=0, l(I,A,B,N).
c2. l(A,B,C,D) :- 2*A+ -1*B>=0,-1*A+1*D>0,-1*A+1*B>=0,3*A+ -1*B+ -1*C=0,
                  1*A+ -1*E= -1,l_body(B,C,F,G),l(E,F,G,D).
c3. l(A,B,C,D):- 3*A+ -3*D>0,1*D>0,2*A+ -1*B>=0,-3*A+3*D> -3,
                  -1*A+1*B>=0,3*A+ -1*B+ -1*C=0.
c4. l(A,B,C,D):- false.
c5. l_body(A,B,C,D) :- -1*A+2*B>=0, 2*A+ -1*B>=0,
                       1*A+ -1*C= -1,1*B+ -1*D= -2.
c6. l_body(A,B,C,D) :- -1*A+2*B>=0,2*A+ -1*B>=0,1*A+ -1*C= -2,1*B+ -1*D= -1.
```

Figure 2: Example program *t4.pl* [6] with strengthened constraints

C^a represented propagation "bottom-up" (via the answers). Furthermore, note that $C^a \to C^q$ since the answers for p are always stronger than the calls to p. Thus it suffices to add the constraint C^a to $\mathcal{B}$.

Specialisation by strengthening the constraints preserves the answers of the goal with respect to which the query-answer transformation was performed. In particular, in our application we have the following property.

Property 2. *If P is a set of CHCs and P_{false} is the set obtained by strengthening the clause constraints as just described, then $P \models \mathsf{false}$ if and only if $P_{\mathsf{false}} \models \mathsf{false}$.*

The result of strengthening the constraints in Figure 1, using the query-answer program with respect to the goal `false`, is shown in Figure 2. Note that the constraint in clause c4 is strengthened to `false`.

4.4 Analysis of the model of the specialised clauses

It may be that the clauses P_{false} do not contain a clause with head false. In this case safety is proven, since clearly $P_{\mathsf{false}} \not\models \mathsf{false}$. If this check fails, the convex polyhedral analysis is now run on the clauses P_{false}. As the experiments later show, safety is often provable by checking the resulting model; if no constrained fact for false is present, then $P_{\mathsf{false}} \not\models \mathsf{false}$. If safety is not proven, there are two possibilities: the approximate model is not precise enough, but *P* has a model, or there is a proof of false. To distinguish these we proceed to try to refine the clauses by splitting predicates.

5 Safety Check and Program Refinement

This section outlines a procedure for safety check, counterexample analysis and refinement. Refinement is considered when a proof of safety or an existence of a real counterexample (that is, a proof of false cannot be established.

Safety check and counterexample analysis The absence of a constrained fact for predicate false in the over-approximation proves that the given set of CHCs is safe. If safety can not be shown, our implementation of the convex polyhedron analysis produces a derivation tree for false as a trace term which we define formally below. For our program in Figure 1, the set of constrained facts representing the approximate model is shown below.

```
f1. l_body(A,B,C,D) :- 1*B+ -1*D>= -2,-1*B+1*D>=1,-1*A+2*B>=0, 2*A+ -1*B>=0,
                       1*A+1*B+ -1*C+ -1*D= -3.
```

```
f2. false :- true.
f3. l(A,B,C,D) :- 1*D>0,2*A+ -1*B>=0,-1*A+1*B>=0,-3*A+3*D> -3,
                  3*A+ -1*B+ -1*C=0.
```

Since there is a constrained fact for false, the shortest derivation for it is found, using clause c1 followed by clause c3. This will be represented as a *trace term* c1(c3), which is formally defined below. The idea of trace terms to capture the shape of derivations was introduced by Gallagher and Lafave [18].

AND-trees and trace terms.   Each CHC is associated with an identifier, as shown in Figure 1. These identifiers are treated as constructors whose arity is the number of non-constraint atoms in the clause body. The following definitions of derivations and trace terms is adapted from [18].

An *AND-tree* is a tree each of whose nodes is labelled by an atom and a clause, such that

1. each non-leaf node is labelled by a clause $A \leftarrow C, A_1, \ldots, A_k$ and an atom A, and has children labelled by $A_1, \ldots, A_k$,

2. each leaf node is labelled by a clause $A \leftarrow C$ and an atom A.

We assume that the variables in node labels are renamed appropriately, details are not given here. Any finite derivation corresponds to an AND-tree, and each AND-tree T can be associated with a trace term $\text{tr}(T)$ defined as:

1. c_j, if T is a single leaf node labelled by the clause of form $A \leftarrow C$ with identifier c_j; or

2. $c_i(\text{tr}(T_1), \ldots, \text{tr}(T_n))$, if T is labelled by the clause with identifier by c_i, and has subtrees $T_1, \ldots, T_n$.

A trace-term uniquely defines an AND-tree (up to renaming of variables). The set of constraints of an AND-tree, represented as $\text{constr}(T)$ is

1. C, if T is a single leaf node labelled by the clause of form $A \leftarrow C$; or

2. $C \cup \bigcup_{i=1..n}(\text{constr}(T_i))$ if T is labelled by the clause $A \leftarrow C, A_1, \ldots, A_k$ and has subtrees $T_1, \ldots, T_n$.

We say that an AND-tree T is satisfiable if $\text{SAT}(\text{constr}(T))$. Let T be an AND-tree whose root is labelled by atom A. Define $\text{proj}(T)$ to be $\text{constr}(T)|_{\text{vars}(A)}$.

Interpolants.   Given two sets of constraints C_1, C_2 such that $C_1 \cup C_2$ is unsatisfiable, a (Craig) interpolant is a constraint I with (1) $C_1 \subseteq I$, (2) $I \cup C_2$ is unsatisfiable and (3) I contains only variables common to C_1 and C_2. We implemented the algorithm from [31] for interpolants for linear constraints.

Given an AND-tree T where $\neg\text{SAT}(\text{constr}(T))$, we can construct an interpolant for each non-root node of T, also known as tree interpolants. Let T' be a sub-tree of T, whose root is labelled with A'. Then the interpolant I associated with A' is defined as above where $C_1 = \text{constr}(T')$ and $C_2 = \text{constr}(T) \setminus C_1$, and the interpolants of subtree of T' together with the constraints at the root of T' implies I. Note that by construction of the AND-tree, the only variables in common between C_1 and C_2 (and hence in I) are the variables in A', the label of T'. More details on tree interpolation can be found in [8].

The set $\text{interpolant}(T)$ is the set of constrained facts $A \leftarrow I$, for all non-root nodes of T labelled by atom A with interpolant I as defined above.

Counterexample checking.   Given a trace term, let T be the corresponding AND-tree. We report that the CHCs have no model if $\text{SAT}(\text{constr}(T))$, and our procedure terminates. For our example it can be verified that $\text{SAT}(\text{constr}(\text{c1}(\text{c3})))$ does not hold, so the trace c1(c3) is a false alarm. We now use the interpolants to split predicates and try to get a more precise approximation of the model.

From the trace term `c1(c3)` in the running example we derive interpolant(`c1(c3)`) = {*I*} where
I = `l(A,B,C,D)` ← A + −3 ∗ B + C + D =< 0.

We then split the constrained facts in the approximation of the model, using the corresponding interpolants and their negations. In the example we split constrained fact `f3` by strengthening its constraint with *I* and ¬*I* respectively. Fioravanti *et al.* use a related technique for splitting clauses [16]. Strengthening first with *I* we get

```
l(A,B,C,D):- D>0,2*A+ -1*B>=0,-1*A+1*B>=0,-3*A+3*D> -3,
             3*A+ -1*B+ -1*C=0,A+ -3*B+C+D=<0
```
which after simplification becomes
```
l(A,B,C,D)  :- -4*A+4*B+ -1*D>=0,1*D>0,-3*A+3*D> -3,2*A+ -1*B>=0,
               3*A+ -1*B+ -1*C=0.
```
We follow the same step with ¬*I* and obtain the following set of constrained facts.
```
l(A,B,C,D)  :- -4*A+4*B+ -1*D>=0,1*D>0,-3*A+3*D> -3,2*A+ -1*B>=0,
               3*A+ -1*B+ -1*C=0.
l(A,B,C,D)  :- 4*A+ -4*B+1*D>0,-1*A+1*B>=0,-3*A+3*D> -3,2*A+ -1*B>=0,
               3*A+ -1*B+ -1*C=0.
```
These together with `f1` and `f2` give us a new set of constrained facts, which forms the input to the refinement phase of our procedure.

Refinement by Predicate Splitting.   Refinement consists of obtaining a specialised set of CHCs from a given set of constrained facts and input set of CHCs. We do this by using polyvariant specialisation (PS) based on the method of multiple specialisation [32] with a property-based abstract domain based on the given set of constrained facts. PS is a program specialisation which introduces several new predicates corresponding to specialised versions of the same predicate. Polyvariant specialisation brings the expressive power of disjunctive predicates into the analysis [17]. Space does not permit a more detailed description. For our running example we obtain a split of the predicate `l` into `l_1` and `l_3` , and the specialised program is as follows.

```
false :-          1*A>0,1*B=0,1*C=0,1*D=0,l_3(B,C,D,A).
l_3(A,B,C,D)  :- 2*A+ -1*B>=0,-1*A+1*B>=0,-1*A+1*D>0,4*A+ -4*B+1*D>0,
                 3*A+ -1*B+ -1*C=0, A+ -1*E= -1,l_body_2(B,C,F,G),l_1(E,F,G,D).
l_3(A,B,C,D)  :- 4*A+ -4*B+1*D>0,3*A+ -3*D>0,-1*A+1*B>=0,-3*A+3*D> -3,
                 3*A+ -1*B+ -1*C=0.
l_1(A,B,C,D)  :- 2*A+ -1*B>=0,-1*A+1*B>=0,-1*A+1*D>0,3*A+ -1*B+ -1*C=0,
                 1*A+ -1*E= -1,l_body_2(B,C,F,G),l_1(E,F,G,D).
l_1(A,B,C,D)  :- 3*A+ -3*D>0,2*A+ -1*B>=0,1*D>0,-1*A+1*B>=0,-3*A+3*D> -3,
                 3*A+ -1*B+ -1*C=0.
l_body_2(A,B,C,D)  :- 2*A+ -1*B>=0,-1*A+2*B>=0,1*A+ -1*C= -1,1*B+ -1*D= -2.
l_body_2(A,B,C,D)  :- 2*A+ -1*B>=0,-1*A+2*B>=0,1*A+ -1*C= -2,1*B+ -1*D= -1.
```

The next iteration continues with this specialised program. The intention of splitting and PS is to guarantee progress of refinement, that is, a counterexample once eliminated never occurs again. Our procedure does not guarantee progress, that is, the same spurious counterexamples might appear in subsequent iterations, but in practice we find the polyvariant specialisation usually eliminates the given counterexample. The large number of constants in the above examples are derived during invariants computation. In the next iteration, our example terminates with a real counter example, thus proving our example program unsafe (over the real numbers).

Toolchain.   Our verification procedure is summarised in Figure 3, which is divided into three parts, an *abstractor* (inside green dotted box), followed by a *safety check* and *counterexample analyser* and *refiner* (inside red box). It should be noted that the tools inside the green and red boxes produce new set of CHCs by specialisation.
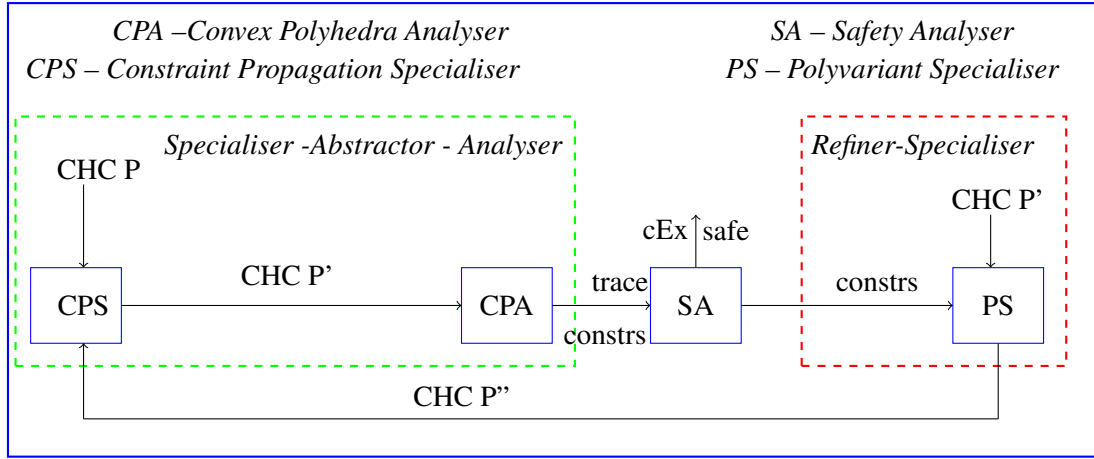

Figure 3: *CHC verification toolchain.*

The effects of CPA and PS in our procedure complement each other and the CPA model gets more accurate during refinement which allows generation of better specialised programs. In essence, it marries the effectiveness of CPA with PS.

6   Experiments

Table 1 presents the results of applying our toolchain depicted in Figure 3 to a number of benchmark programs taken from the repository of Horn clause benchmarks in SMT-LIB2[2] and other sources including [19, 26, 22, 5, 14]. The experiments were carried out using a computer, Intel(R) X5355 having 4 processors (each @ 2.66GHz) and total memory of 6 GB. Debian 5 (64 bit) is the Operating System running in it and we set 2 minutes of timeout for each experiment. Our tool-chain is implemented in 32-bit Ciao Prolog [9][3] and the Parma Polyhedra Library [1][4] for this purpose.

In Table 1, columns Program, "n" , Result and time (sec) respectively represent the benchmark program, the number of refinement iterations necessary to verify a given property, the results of verification and the time (in seconds) to verify them. Value 0 in column "n" means that no refinement is necessary, whereas value greater than 0 indicates the actual number of iterations necessary and value "-" means that these programs are beyond the reach of our current tool within the given time limit. Problems marked with (*) were not handled by our tool-chain since their solution generates numbers which do not fit in 32 bits, the limit of our Ciao Prolog implementation. Problems such as systemc-token-ring.01-safeil.c contain complicated loop structure with large strongly connected components in the predicate dependency graph and our convex polyhedron analysis tool is unable to derive the required invariant.

---

[2]https://svn.sosy-lab.org/software/sv-benchmarks/trunk/clauses/

[3]http://ciao-lang.org/

[4]http://bugseng.com/products/ppl/

Program	n	Result	time (secs)	Program	n	Result	time (secs)
MAP-disj.c.map.pl	0	safe	1.0	jaffex1c.pl	0	safe	0.01
MAP-forward.c.map.pl	0	safe	1.0	jaffex1a.pl	0	safe	0.01
t1.pl	0	safe	0.01	qrdcmp.smt2	0	safe	118.0
t1-a.pl	0	safe	0.01	choldc.smt2	0	safe	19.0
t2.pl	0	safe	0.01	lop.smt2	0	safe	39.0
t3.pl	0	safe	1.0	pzextr.smt2	0	safe	40.0
t4.pl	1	unsafe	1.0	qrsolv.smt2	0	safe	18.0
t5.pl	0	safe	0.01	tridag.smt2	0	safe	13.0
MAP-disj.c-scaled.pl	0	safe	1.0	systemc-pc-sfifo_1	0	unsafe	12.0
INVGEN-id-build	0	safe	1.0	loops-terminator	0	unsafe	0.01
INVGEN-nested5	0	safe	1.0	loops-for-bounded	3	unsafe	5.0
INVGEN-nested6	0	safe	117.0	TRACER-testabs15	0	safe	1.0
INVGEN-nested8	0	safe	1.0	INVGEN-apache-esc-abs	0	safe	2.0
INVGEN-svdsomeloop	0	safe	3.0	DAGGER-barbr.map.c	0	safe	119.0
INVGEN-svd1	2	safe	13.0	systemc-token-ring.01-safeil.c	-	?	-
INVGEN-svd4	0	safe	5.0	sshs3-srvr1a-safeil.c(*)	-	?	-
loops-count-up-down	0	unsafe	1.0	sshs3-srvr1b-safeil.c	-	?	-
loops-sum04	8	unsafe	2.0	amebsa.smt2	-	?	-
dfpp12.pl	0	safe	0.01	bandec.smt2(*)	-	?	-
TRACER-testloop27	1	unsafe	1.0	TRACER-testloop28	-	?	-
TRACER-testloop8	0	unsafe	0.01	crank.smt2	-	?	-
jaffex1b.pl	0	safe	0.01	pldi12.pl	-	?	-
jaffex1d.pl	0	safe	0.01	loops-sum01	-	?	-

Table 1: Experimental results on CHC benchmark problems

The results of our procedure in a larger set of benchmarks obtained from previous sources are summarised in Table 2. Though our tool-chain is not optimized at all, the overall result shows that it compares favourably with other advanced verification tools like HSF [20], VeriMAP [14], TRACER [25] etc. in both time and the number of problems solved, and thus showing the effectiveness of our approach.

	without refinemet	with refinement
solved (safe/unsafe)	160 (142/18)	181 (158/23)
unknown/ timeout	49/7	-/35
total time	1293	3410
average time (secs)	5.98	18.73

Table 2: Experimental results on 216 CHC verification problems, where "-" means not relevant.

7   Related Work

Verification of CLP programs using abstract interpretation and specialisation has been studied for some time. The use of an over-approximation of the semantics of a program can be used to establish safety properties – if a state or property does not appear in an over-approximation, it certainly does not appear in the actual program behaviour. A general framework for logic program verification through abstraction was described by Levi [29].

The use of program transformation to verify properties of logic programs was pioneered by Pettorossi and Proietti [30] and Leuschel [28]. Transformations that preserve the minimal model (or other suitable models) of logic programs are applied systematically to make properties explicit. For example, if a program can be transformed to one containing a clause $A \leftarrow true$ then A is a consequence of the program.

Recent work by De Angelis *et al.* [13, 14] applies a specialisation approach to the Horn clause verification problem as discussed here, namely, with integrity constraints expressing the properties to be proved. Both our approach and theirs repeatedly apply specialisations preserving the property to be proved. However the difference is that their specialisation techniques are based on unfold-fold transformations, with a sophisticated control procedure controlling unfolding and generalisation. Our specialisations are restricted to strengthening of constraints or polyvariant splitting based on local conditions. Their test for success or failure is a simple syntactic check, whereas ours is based on an abstract interpretation to derive an over-approximation.

Counterexample guided abstraction refinement (CEGAR) [10] has been successfully used in verification to automatically refine (predicate) abstractions to reduce false alarms but not much has been explored in refining abstractions in the convex polyhedral domain. See [7, 21] for more details about the use of interpolation in refinement. A number of tools implementing predicate abstraction and refinement are available, such as HSF [20] and BLAST [2]. TRACER [19] is a verification tool based on CLP that uses symbolic execution.

Informally one can say that approaches differ in where the "hard work" is performed. In the work of De Angelis *et al.* the specialisation procedure is the core, whereas in the CEGAR approaches the refinement step is crucial, and interpolation plays a central role. In our approach, by contrast, most of the hard work is done by the abstract interpretation, which finds useful invariants as well as propagating constraints globally. The main problem is to find effective ways of refining polyhedral abstractions. Finding the most effective balance between specialisation, abstraction and refinement techniques is a matter of ongoing research.

8   Conclusion and Future works

We described an iterative procedure for Horn clause verification which interleaves abstract interpretation with specialisation. A specialised set of CHCs is produced first by strengthening the constraints in the given clauses using the results of the abstract interpretation. Then the procedure terminates if an abstract interpretation of the resulting program is sufficient to verify the required properties, otherwise, a polyvariant specialisation guided by an abstract counterexample is performed using the inferred constraints as well as interpolated constraints.

In the future, we would like to find a way of ensuring progress of refinement, maybe using the powerset polyhedra domain, and also interface our toolchain with SMT solvers for satisfiability checking and interpolant generation.

References

[1] R. Bagnara, P. M. Hill & E. Zaffanella (2008): *The Parma Polyhedra Library: Toward a Complete Set of Numerical Abstractions for the Analysis and Verification of Hardware and Software Systems*. Science of Computer Programming 72(1–2), pp. 3–21. Available at `http://dx.doi.org/10.1016/j.scico.2007.08.001`.

[2] T. Ball, V. Levin & S. K. Rajamani (2011): *A decade of software model checking with SLAM*. Commun. ACM 54(7), pp. 68–76. Available at `http://doi.acm.org/10.1145/1965724.1965743`.

[3] F. Bancilhon, D. Maier, Y. Sagiv & J. Ullman (1986): *Magic Sets and other strange ways to implement logic programs*. In: *Proceedings of the 5^{th} ACM SIGMOD-SIGACT Symposium on Principles of Database Systems*. Available at `http://dx.doi.org/10.1145/6012.15399`.

[4] F. Benoy & A. King (1996): *Inferring Argument Size Relationships with CLP(R)*. In J. P. Gallagher, editor: *Logic-Based Program Synthesis and Transformation (LOPSTR'96)*, Springer-Verlag Lecture Notes in Computer Science 1207, pp. 204–223. Available at `http://dx.doi.org/10.1007/3-540-62718-9_12`.

[5] D. Beyer (2013): *Second Competition on Software Verification - (Summary of SV-COMP 2013)*. In N. Piterman & S. A. Smolka, editors: *TACAS*, Lecture Notes in Computer Science 7795, Springer, pp. 594–609. Available at `http://dx.doi.org/10.1007/978-3-642-36742-7_43`.

[6] D. Beyer, T. A. Henzinger, R. Majumdar & A. Rybalchenko (2007): *Path invariants*. In J. Ferrante & K. S. McKinley, editors: *PLDI*, ACM, pp. 300–309. Available at `http://doi.acm.org/10.1145/1250734.1250769`.

[7] N. Bjørner, K. L. McMillan & A. Rybalchenko (2013): *On Solving Universally Quantified Horn Clauses*. In F. Logozzo & M. Fähndrich, editors: *SAS*, Lecture Notes in Computer Science 7935, Springer, pp. 105–125. Available at `http://dx.doi.org/10.1007/978-3-642-38856-9_8`.

[8] R. Blanc, A. Gupta, L. Kovács & B. Kragl (2013): *Tree Interpolation in Vampire*. In K. L. McMillan, A. Middeldorp & A. Voronkov, editors: *LPAR*, Lecture Notes in Computer Science 8312, Springer, pp. 173–181. Available at `http://dx.doi.org/10.1007/978-3-642-45221-5_13`.

[9] F. Bueno, D. Cabeza, M. Carro, M. Hermenegildo, P. López-García & G. Puebla (1997): *The Ciao Prolog system. Reference manual*. Technical Report CLIP3/97.1, School of Computer Science, Technical University of Madrid (UPM). Available from http://www.clip.dia.fi.upm.es/.

[10] E. M. Clarke, O. Grumberg, S. Jha, Y. Lu & H. Veith (2003): *Counterexample-guided abstraction refinement for symbolic model checking*. J. ACM 50(5), pp. 752–794. Available at `http://doi.acm.org/10.1145/876638.876643`.

[11] P. Cousot & R. Cousot (1977): *Abstract Interpretation: A Unified Lattice Model for Static Analysis of Programs by Construction or Approximation of Fixpoints*. In R. M. Graham, M. A. Harrison & R. Sethi, editors: *Conference Record of the Fourth ACM Symposium on Principles of Programming Languages, Los Angeles, California, USA, January 1977*, ACM, pp. 238–252. Available at `http://dl.acm.org/citation.cfm?id=512950`.

[12] P. Cousot & N. Halbwachs (1978): *Automatic Discovery of Linear Restraints Among Variables of a Program*. In A. V. Aho, S. N. Zilles & T. G. Szymanski, editors: *Conference Record of the Fifth Annual ACM Symposium on Principles of Programming Languages, Tucson, Arizona, USA, January 1978*, ACM Press, pp. 84–96. Available at `http://dl.acm.org/citation.cfm?id=512760`.

[13] E. De Angelis, F. Fioravanti, A. Pettorossi & M. Proietti (2014): *Program verification via iterated specialization*. Sci. Comput. Program. 95, pp. 149–175. Available at `http://dx.doi.org/10.1016/j.scico.2014.05.017`.

[14] E. De Angelis, F. Fioravanti, A. Pettorossi & M. Proietti (2014): *VeriMAP: A Tool for Verifying Programs through Transformations*. In E. Ábrahám & K. Havelund, editors: *TACAS*, Lecture Notes in Computer Science 8413, Springer, pp. 568–574. Available at `http://dx.doi.org/10.1007/978-3-642-54862-8_47`.

[15] S. Debray & R. Ramakrishnan (1994): *Abstract Interpretation of Logic Programs Using Magic Transformations*. Journal of Logic Programming 18, pp. 149–176. Available at `http://dx.doi.org/10.1016/0743-1066(94)90050-7`.

[16] F. Fioravanti, A. Pettorossi & M. Proietti (2002): *Specialization with Clause Splitting for Deriving Deterministic Constraint Logic Programs*. In: *In Proc. IEEE Conference on Systems, Man and Cybernetics, Hammamet*, IEEE Press. Available at `http://dx.doi.org/10.1109/ICSMC.2002.1167971`.

[17] F. Fioravanti, A. Pettorossi, M. Proietti & V. Senni (2013): *Controlling Polyvariance for Specialization-based Verification*. Fundam. Inform. 124(4), pp. 483–502. Available at `http://dx.doi.org/10.3233/FI-2013-845`.

[18] J. P. Gallagher & L. Lafave (1996): *Regular Approximation of Computation Paths in Logic and Functional Languages*. In O. Danvy, R. Glück & P. Thiemann, editors: *Partial Evaluation*, Springer-Verlag Lecture Notes in Computer Science 1110, pp. 115–136. Available at `http://dx.doi.org/10.1007/3-540-61580-6_7`.

[19] G. Gange, J. A. Navas, P. Schachte, H. Søndergaard & P. J. Stuckey (2013): *Failure tabled constraint logic programming by interpolation*. TPLP 13(4-5), pp. 593–607. Available at `http://dx.doi.org/10.1017/S1471068413000379`.

[20] S. Grebenshchikov, A. Gupta, N. P. Lopes, C. Popeea & A. Rybalchenko (2012): *HSF(C): A Software Verifier Based on Horn Clauses - (Competition Contribution)*. In C. Flanagan & B. König, editors: *TACAS*, LNCS 7214, Springer, pp. 549–551. Available at `http://dx.doi.org/10.1007/978-3-642-28756-5_46`.

[21] A. Gupta, C. Popeea & A. Rybalchenko (2011): *Solving Recursion-Free Horn Clauses over LI+UIF*. In H. Yang, editor: *APLAS*, Lecture Notes in Computer Science 7078, Springer, pp. 188–203. Available at `http://dx.doi.org/10.1007/978-3-642-25318-8_16`.

[22] A. Gupta & A. Rybalchenko (2009): *InvGen: An Efficient Invariant Generator*. In A. Bouajjani & O. Maler, editors: *CAV*, Lecture Notes in Computer Science 5643, Springer, pp. 634–640. Available at `http://dx.doi.org/10.1007/978-3-642-02658-4_48`.

[23] N. Halbwachs, Y. E. Proy & P. Raymound (1994): *Verification of Linear hybrid systems by means of convex approximations*. In: *Proceedings of the First Symposium on Static Analysis*, LNCS 864, pp. 223–237. Available at `http://dx.doi.org/10.1007/3-540-58485-4_43`.

[24] J. Jaffar & M. Maher (1994): *Constraint Logic Programming: A Survey*. Journal of Logic Programming 19/20, pp. 503–581. Available at `http://dx.doi.org/10.1016/0743-1066(94)90033-7`.

[25] J. Jaffar, V. Murali, J. A. Navas & A. E. Santosa (2012): *TRACER: A Symbolic Execution Tool for Verification*. In P. Madhusudan & S. A. Seshia, editors: *CAV*, Lecture Notes in Computer Science 7358, Springer, pp. 758–766. Available at `http://dx.doi.org/10.1007/978-3-642-31424-7_61`.

[26] J. Jaffar, J. A. Navas & A. E. Santosa (2011): *Unbounded Symbolic Execution for Program Verification*. In S. Khurshid & K. Sen, editors: *RV*, Lecture Notes in Computer Science 7186, Springer, pp. 396–411. Available at `http://dx.doi.org/10.1007/978-3-642-29860-8_32`.

[27] L. Lakhdar-Chaouch, B. Jeannet & A. Girault (2011): *Widening with Thresholds for Programs with Complex Control Graphs*. In T. Bultan & P.-A. Hsiung, editors: *ATVA 2011*, Lecture Notes in Computer Science 6996, Springer, pp. 492–502. Available at `http://dx.doi.org/10.1007/978-3-642-24372-1_38`.

[28] M. Leuschel & T. Massart (1999): *Infinite State Model Checking by Abstract Interpretation and Program Specialisation*. In A. Bossi, editor: *LOPSTR'99*, Lecture Notes in Computer Science 1817, Springer, pp. 62–81. Available at `http://dx.doi.org/10.1007/10720327_5`.

[29] G. Levi (2000): *Abstract Interpretation Based Verification of Logic Programs*. Electr. Notes Theor. Comput. Sci. 40, p. 243. Available at `http://dx.doi.org/10.1016/S1571-0661(05)80052-0`.

[30] A. Pettorossi & M. Proietti (2000): *Perfect Model Checking via Unfold/Fold Transformations*. In J. W. Lloyd, V. Dahl, U. Furbach, M. Kerber, K.-K. Lau, C. Palamidessi, L. M. Pereira, Y. Sagiv & P. J. Stuckey, editors: *Computational Logic*, Lecture Notes in Computer Science 1861, Springer, pp. 613–628. Available at `http://dx.doi.org/10.1007/3-540-44957-4_41`.

[31] A. Rybalchenko & V. Sofronie-Stokkermans (2010): *Constraint solving for interpolation. J. Symb. Comput.* 45(11), pp. 1212–1233. Available at `http://dx.doi.org/10.1016/j.jsc.2010.06.005`.

[32] W. H. Winsborough (1989): *Path-Dependent Reachability Analysis for Multiple Specialization.* In E. L. Lusk & R. A. Overbeek, editors: *NACLP*, MIT Press, pp. 133–153. Available at `http://dblp.uni-trier.de/db/conf/slp/slp89.html#Winsborough89`.