

Solving non-linear Horn clauses using a linear Horn clause solver *

Bishoksan Kafle

Roskilde University, Denmark

kafle@ruc.dk

John P. Gallagher

Roskilde University, Denmark

IMDEA Software Institute, Spain

jpg@ruc.dk

Pierre Ganty

IMDEA Software Institute, Spain

pierre.ganty@imdea.org

In this paper we show that checking satisfiability of a set of non-linear Horn clauses (also called a non-linear Horn clause program) can be achieved using a solver for linear Horn clauses. We achieve this by interleaving a *program transformation* with a satisfiability checker for linear Horn clauses (also called a *solver* for linear Horn clauses). The program transformation is based on the notion of *tree dimension*, which we apply to a set of non-linear clauses, yielding a set whose derivation trees have bounded dimension. Such a set of clauses can be linearised. The main algorithm then proceeds by applying the linearisation transformation and solver for linear Horn clauses to a sequence of sets of clauses with successively increasing dimension bound. The approach is then further developed by using a solution of clauses of lower dimension to (partially) linearise clauses of higher dimension. We constructed a prototype implementation of this approach and performed some experiments on a set of verification problems, which shows some promise.

1 Introduction

Many software verification problems can be reduced to checking satisfiability of a set of Horn clauses (the *verification conditions*). In this paper we propose an approach for checking satisfiability of a set of non-linear Horn clauses (clauses whose body contains more than one non-constraint atom) using a linear Horn clause solver. A *program transformation* based on the notion of *tree dimension* is applied to a set of non-linear Horn clauses; this gives a set of clauses that can be linearised and then solved using a *linear solver* for Horn clauses. This combination of dimension-bounding, linearisation and then solving with a linear solver is repeated for successively increasing dimension. The dimension of a tree is a measure of its non-linearity – for example a linear tree (whose nodes have at most one child) has dimension zero while a complete binary tree has dimension equal to its height.

A given set of Horn clauses P can be transformed into a new set of clauses $P^{\leq k}$, whose derivation trees are the subset of P 's derivation trees with dimension at most k . It is known that $P^{\leq k}$ can be transformed to a linear set of clauses preserving satisfiability; hence if we can find a model of the linear set of clauses then the original clauses $P^{\leq k}$ also have a model.

The algorithm terminates with success if a model (solution) M of $P^{\leq k}$ is also a model (after appropriate translation of predicate names) of P . However if M is not a solution of P , then we proceed to generate $P^{\leq k+1}$ and repeat the procedure. The algorithm terminates with failure if $P^{\leq k}$ is shown to be unsatisfiable for some k , since this implies that P is also unsatisfiable.

*The research leading to these results has been supported by EU FP7 project 318337, *ENTRA - Whole-Systems Energy Transparency*, EU FP7 project 611004, *coordination and support action ICT-Energy*, EU FP7 project 610686, *POLCA - Programming Large Scale Heterogeneous Infrastructures*, Madrid Regional Government project S2013/ICE-2731, *N-Greens Software - Next-Generation Energy-Efficient Secure Software*, and the Spanish Ministry of Economy and Competitiveness project No. TIN2015-71819-P, *RISCO - Rigorous analysis of Sophisticated COncurrent and distributed systems*.

```

c1. fib(A, B):- A>=0, A=<1, B=A.
c2. fib(A, B) :- A > 1, A2 = A - 2, fib(A2, B2),
      A1 = A - 1, fib(A1, B1), B = B1 + B2.
c3. false:- A>5, fib(A,B), B<A.

```

Figure 1: Example CHCs Fib defining a Fibonacci function.

A more sophisticated version of the algorithm attempts to use the model M of $P^{\leq k}$ to (partially) linearise $P^{\leq k+1}$. We can exploit the model of $P^{\leq k}$ in the following way; if $P^{\leq k+1}$ has a counterexample that does not use the (approximate) solution M for $P^{\leq k}$, then P is unsatisfiable. We continue this process successively for increasing value of k until we find a solution or a counterexample to P , or until resources are exhausted.

As an example program, we consider a set of constrained Horn clauses P in Figure 1 which defines the Fibonacci function. It is an interesting problem since its derivations are trees, whose dimensions depend on an input argument. The last clause represents a property of the Fibonacci function expressed as an integrity constraint.

We have made a prototype implementation of this approach and performed some experiments on a set of software verification problems, which shows some promise. The main contributions of this paper are as follows.

1. We present a linearisation procedure for dimension-bounded Horn clauses using partial evaluation (Section 3).
2. We give an algorithm for solving a set of non-linear Horn clauses using a linear Horn clause solver (Section 4).
3. We demonstrate the feasibility of our approach in practice applying it to non-linear Horn clause problems (Section 5).

2 Preliminaries

A constrained Horn clause (CHC) is a first order formula of the form $p(X) \leftarrow \mathcal{C}, p_1(X_1), \dots, p_k(X_k)$ ($k \geq 0$) (using Constraint Logic Programming syntax), where \mathcal{C} is a conjunction of *constraints* with respect to some constraint theory, X_i, X are (possibly empty) vectors of distinct *variables*, p_1, \dots, p_k, p are *predicate symbols*, $p(X)$ is the *head* of the clause and $\mathcal{C}, p_1(X_1), \dots, p_k(X_k)$ is the *body*. An atomic formula, or simply *atom*, is a formula $p(t)$ where p is a non-constraint predicate symbol and t a tuple of arguments. Atoms are sometimes written as A, B or H , possibly with sub- or superscripts.

A clause is called *non-linear* if it contains more than one atom in the body, otherwise it is called *linear*. A set of Horn clauses P is called linear if P only contains linear clauses, otherwise it is called non-linear. *Integrity constraints* are a special kind of Horn clauses whose head is *false* where *false* is always interpreted as FALSE. A set of Horn clauses is sometimes called a (*constraint logic*) *program*.

An *interpretation* of a set of CHCs is represented as a set of *constrained facts* of the form $A \leftarrow \mathcal{C}$ where A is an atomic formula $p(Z)$ where Z is a tuple of distinct variables and \mathcal{C} is a constraint over Z with respect to some constraint theory. An interpretation that makes each clause in P TRUE is called a *model* of P . We say a set of Horn clause P (including integrity constraints) is safe (solvable) iff it has a model. In some works e.g. [5, 29], a model is also called a *solution* and we use them interchangeably in this paper.

A labeled tree $c(t_1, \dots, t_k)$ ($n \geq 0$) is a tree whose nodes are labeled by identifiers, where c is the label of the root and t_1, \dots, t_k are labeled trees, the children of the root.

Definition 1 (Tree dimension (adapted from [12])) Given a labeled tree $t = c(t_1, \dots, t_k)$, the tree dimension of t represented as $\dim(t)$ is defined as follows:

$$\dim(t) = \begin{cases} 0 & \text{if } k = 0 \\ \max_{i \in [1..k]} \dim(t_i) & \text{if there is a unique maximum} \\ \max_{i \in [1..k]} \dim(t_i) + 1 & \text{otherwise} \end{cases}$$

Given a set of Horn clauses, we associate with each clause $p(X) \leftarrow \mathcal{C}, p_1(X_1), \dots, p_k(X_k)$ a unique identifier c whose *arity* is k . A labeled tree whose node labels are clause identifiers is called a *trace tree*. The term $c(t_1, \dots, t_k)$ is a *trace tree* which represents a derivation tree whose root is the head of the clause identified by c , and t_1, \dots, t_k represent the derivation trees for the atomic formulas in the body (we assume that variables are renamed appropriately). There is a one-one correspondence between *trace trees* and derivation trees of Horn clauses up to variable renaming. Thus when we speak about the dimension of a Horn clause derivation, we refer to the dimension of its corresponding *trace tree*.

Using the clauses shown in Figure 1 along with their identifiers, Figure 2 (a) shows a *trace tree* $t = c_3(c_2(c_2(c_1, c_1), c_1))$ and Figure 2 (b) shows its tree dimension. It can be seen that $\dim(t) = 1$.




Figure 2: (a) a *trace tree* and (b) its tree dimension.

To make the paper self contained, we describe the transformation to produce a dimension-bounded set of clauses. Given a set of CHCs P and $k \in \mathbb{N}$, we split each predicate p occurring in P into the predicates $p^{\leq d}$ and $p^{=d}$ where $d \in \{0, 1, \dots, k\}$. An atom with predicate $p^{\leq d}$ or $p^{=d}$ is denoted $H^{\leq d}$ or $H^{=d}$ respectively. Such atoms have derivation trees of dimension at most d and exactly d respectively.

Definition 2 (At-most- k -dimension program $P^{\leq k}$) Let P be a set of CHCs. $P^{\leq k}$ consists of the following clauses (adapted from [28, 24]):

1. *Linear clauses:*

If $H \leftarrow \mathcal{C} \in P$, then $H^{=0} \leftarrow \mathcal{C} \in P^{\leq k}$.

If $H \leftarrow \mathcal{C}, B_1 \in P$ then $H^{=d} \leftarrow \mathcal{C}, B_1^d \in P^{\leq k}$ for $0 \leq d \leq k$.

2. *Non-linear clauses:*

If $H \leftarrow \mathcal{C}, B_1, B_2, \dots, B_r \in P$ with $r > 1$ and one of the following holds:

- For $1 \leq d \leq k$, and $1 \leq j \leq r$:

Set $Z_j = B_j^{=d}$ and $Z_i = B_i^{\leq d-1}$ for $1 \leq i \leq r \wedge i \neq j$. Then: $H^{=d} \leftarrow \mathcal{C}, Z_1, \dots, Z_r \in P^{\leq k}$.

```

%linear clauses
1. fib(0) (A,B) :- A>=0, A<1, B=A.
2. false(0) :- A>5, B<A, fib(0) (A,B).
%epsilon-clauses
3. false[0] :- false(0).
4. fib[0] (A,B) :- fib(0) (A,B).

```

Figure 3: $\text{Fib}^{\leq 0}$: at-most-0-dimension program of Fib.

- For $1 \leq d \leq k$, and $J \subseteq \{1, \dots, r\}$ with $|J| = 2$:
Set $Z_i = B_i^{\leq d-1}$ if $i \in J$ and $Z_i = B_i^{\leq d-2}$ if $i \in \{1, \dots, r\} \setminus J$. If all Z_i are defined, i.e., $d \geq 2$ if $r > 2$, then: $H^{\leq d} \leftarrow \mathcal{C}, Z_1, \dots, Z_r \in P^{\leq k}$.

3. ε -clauses:

$H^{\leq d} \leftarrow H^{\leq e} \in P^{\leq k}$ for $0 \leq d \leq k$, and every $0 \leq e \leq d \leq k$.

$P^{\leq k}$ is also called the k -dimension-bounded program corresponding to P . When the value of k is not important, any program generated using the Definition 2 is called a dimension-bounded program. The relation between P and its k -dimensional program is given by in the Proposition 1 where \models is the usual “logical consequence” operator.

Proposition 1 (Relation between P and $P^{\leq k}$) Let P be a program and $P^{\leq k}$ ($k \geq 0$) be its k -dimensional program. Let $p(t)$ be an atom where p is a predicate of P and $p^*(t)$ ($\star \in \{=, \leq\}$) be an atom where p^* is a predicate of $P^{\leq k}$. Then we have: $P^{\leq k} \models p^*(t) \implies P \models p(t)$.

In other words, Proposition 1 says that the set of facts that can be derived from $P^{\leq k}$ is a subset of the set of facts that can be derived from P , taking the predicate renaming into account. In this sense $P^{\leq k}$ is an *under-approximation* of P . In particular, if $P^{\leq k} \models \text{false}^*$ then $P \models \text{false}$.

Let S be an interpretation of a dimension-bounded set of clauses $P^{\leq k}$. That is, S is a set of constrained facts of the form $H^{\leq d} \leftarrow \mathcal{C}$ or $H^{\leq d} \leftarrow \mathcal{C}$. An interpretation of P is constructed from S as follows.

Definition 3 ($S^{\uparrow P^{\leq k}}$: an interpretation of P constructed from an interpretation of $P^{\leq k}$) Let S be an interpretation of $P^{\leq k}$. Then $S^{\uparrow P^{\leq k}}$ is the following set of constrained facts.

$$S^{\uparrow P^{\leq k}} = \{p(X) \leftarrow \bigvee \{\mathcal{C} \mid p^{\leq d}(X) \leftarrow \mathcal{C} \in S \vee p^{\leq d}(X) \leftarrow \mathcal{C} \in S\} \mid d \in \{0 \dots k\} \text{ and } p \text{ is a predicate in } P\}$$

The set $S^{\uparrow P^{\leq k}}$ is a disjunctive interpretation of P where the interpretation of p is the disjunction of the interpretations of the corresponding dimension-bounded versions of p in $P^{\leq k}$.

The at-most-0-dimension program of Fib in Figure 1 is depicted in Figure 3. In textual form we represent a predicate $p^{\leq k}$ by $p[k]$ and a predicate $p^{\leq k}$ by $p(k)$. The at-most-1-dimension program of Fib in Figure 1 is depicted in Figure 4. Note that 0-dimension program is included in 1-dimension program. In general, all the clauses in $P^{\leq k}$ are also in $P^{\leq k+1}$. This provides a basis for an iterative strategy for a bounded set of Horn clauses. Since some programs have derivation trees of unbounded dimension, trying to verify a property for its increasing dimension separately is not a practical strategy. It only becomes a viable approach if a solution of $p^{\leq k}$ for some $k \geq 0$ is general enough to hold for all dimensions of P .

```

fib(0)(A,B) :- B=A, A=<1, A>=0.
fib(1)(A,B) :- B=F+D, C=A-2,
               E=A-1, A>1, fib[0](E,F), fib(1)(C,D).
fib(1)(A,B) :- B=F+D, C=A-2, E=A-1,
               A>1, fib[0](C,D), fib(1)(E,F).
fib(1)(A,B) :- B=F+D, C=A-2, E=A-1,
               A>1, fib(0)(C,D), fib(0)(E,F).
false(1) :- B<A, A>5, fib(1)(A,B).
false(0) :- B<A, A>5, fib(0)(A,B).
false[1] :- false(1).
false[1] :- false(0).
false[0] :- false(0).
fib[1](A,B) :- fib(1)(A,B).
fib[1](A,B) :- fib(0)(A,B).
fib[0](A,B) :- fib(0)(A,B).

```

Figure 4: $\text{Fib}^{\leq 1}$: at-most-1-dimension program of Fib.

3 Linearisation strategies for dimension-bounded set of Horn clauses

In this section, we present linearisation strategies for set of clauses of bounded dimension. It is known [1] that a dimension-bounded set of clauses can be linearised, preserving satisfiability. In this section we describe a practical technique for linearisation, based on partial evaluation of an interpreter.

3.1 Linearisation based on partial evaluation

Partial evaluation (PE) has been studied for a variety of languages including logic programs [21, 14, 25, 22, 27]. We follow the pattern of transforming a program (a set of Horn clauses) by specialising an interpreter for that program [13, 22]. Let PE be a partial evaluator, I an interpreter and P an object program. Then the partial evaluation of I with respect to P , denoted $PE(I, P)$, represents the “compilation” of P using the semantics given by I .

We first write an interpreter for Horn clause programs, which is also written as a set of Horn clauses. Given a (possibly empty) conjunction of atoms (called a *goal*) the interpreter constructs a derivation, implementing a standard left-to-right, depth-first search. In the interpreter predicate `solve(Gs)`, Gs is the goal, represented as a list of atoms. The basic step of the interpreter is represented by the clauses for `solve(Gs)` shown in Figure 5. If the conjunction is not empty, its first atom G is selected along with a matching Horn clause $G \leftarrow Cs, B$ in the program being interpreted, where Cs is a conjunction of constraints and B is a conjunction of atoms. This clause is represented by `hornClause(G, Cs, B)` in the interpreter. The body of the clause is conjoined with the remaining goal atoms, and the derivation continues with the new goal $Gs1$. If the conjunction is empty, the derivation is successful (second clause).

To interpret a dimension-bounded set of clauses (say the bound is k), we use the fact that in all successful runs of the interpreter in which goals are selected in increasing order of dimension, the size of the conjunction of goals (that is, the length of the argument of `solve`) has an upper bound related to k . This bound is known as the *index* of the set of clauses and is given as $(i - 1) * k + 1$, where i is the maximum number of non-constraint atoms in the body of clauses [11]. Given this index, we can augment the interpreter with a check on the size of the conjunction, ensuring that it never exceeds the

```

solve([G|Gs]) :-
    hornClause(G,Cs,B), solveConstraints(Cs), append(B,Gs,Gs1),
    solve(Gs1).
solve([]).

```

Figure 5: Depth-first interpreter for Horn clauses

index. In addition, due to the requirement of increasing dimension in the selection of atoms, a left-to-right computation rule is not sufficient; therefore we permute the set of atoms in each clause body, since in at least one permutation the goals will be ordered by dimension. With these changes the interpreter remains complete for clauses of the given maximum index, at the possible cost of some redundancy in the search.

These additions result in the interpreter whose top level is shown in Figure 6. Let the interpreter predicate `solve(Gs, Index, L)` mean that the conjunction of goals `Gs` is to be solved, and `L`, `Index` are numbers representing the size of `Gs` and the maximum size of the stack of goals.

```

go(Index) :-
    solve([false], Index, 1).
solve([G|Gs], Index, L) :-
    hornClause(G,Cs,B), solveConstraints(Cs),
    length(B,L1), L2 is L1+L-1, L2 =< Index,
    perm(B,B1), append(B1,Gs,Gs1),
    solve(Gs1, Index, L2).
solve([], _, _).

```

Figure 6: Interpreter for linearisation

Partial evaluation of the interpreter. Given a set of facts of the form `hornClause(G,Cs,B)` representing the Horn clauses to be linearised, and some value of the index K , the interpreter can be partially evaluated. We use Logen [26] to perform the partial evaluation with respect to a call to `go(Index)`, which initiates a proof of the goal `false` (see first clause of interpreter). All interpreter computations are partially evaluated except for the calls to `solve(Gs, Index, L)` and the execution of constraints within the goal `solveConstraints(Cs)`. Furthermore Logen performs standard structure-flattening and predicate renaming operations, yielding a set of clauses of the form `solve_i(X) :- Cs, solve_j(Y)`, where `solve_i(X)` and `solve_j(Y)` are instantiations of `solve(Gs, Index, L)` and `Cs` is a constraint. Thus the resulting clauses are linear, and furthermore preserve the meaning of the original clauses as given by the interpreter, by correctness of the partial evaluation procedure. The linearisation procedure is independent of the constraint theory underlying the clauses.

Proposition 2 *Let P be a program and $P^{\leq k}$ ($k \geq 0$) be its k -dimensional program. Let i be the maximum number of atoms in clause bodies of $P^{\leq k}$. Let $\text{Index} = (i - 1) * k + 1$. Let P' be a partial evaluation of the interpreter in Figure 6, with respect to P and the goal `go(Index)`. Then $P^{\leq k} \models \text{false}^{\leq k}$ iff $P' \models \text{go}(\text{Index})$.*

Furthermore P' is linear if the partial evaluator follows the strategy described above. Combining Propositions 2 and 1, we conclude that $P' \models \text{go}(\text{Index}) \Rightarrow P \models \text{false}$.

Note that linearisation required partial evaluation of the perm predicate, giving a blow-up in program size related to the length of the clause bodies. This is further discussed at the end of Section 5.

3.2 Obtaining linear over-approximations with a partial model

First we note that the set of predicates in $P^{\leq k}$ is a subset of the set of predicates in $P^{\leq k+1}$. Given a model M for the predicates in $P^{\leq k}$, $P^{\leq k+1}$ can be linearised if we replace each occurrence of a predicate from $P^{\leq k}$ in the body of a clause in $P^{\leq k+1}$ with the corresponding constraint from the model M . The resulting set of clauses is linear since $P^{\leq k+1}$ contains at most one predicate in its body from $P^{\leq k+1}$ which is not in $P^{\leq k}$. Furthermore if $P^{\leq k+1}$ has a model then so does the set of clauses resulting from the replacement; the converse is however not the case since the model M represents an over-approximation of $P^{\leq k}$.

More generally, we can replace any subset of the occurrences of predicates from $P^{\leq k}$ in $P^{\leq k+1}$. We summarise this in the following lemma.

Lemma 1 (Linear over-approximation) *Let M be a model of the predicates in $P^{\leq k}$, represented by a set of “constrained facts” $p(X) \leftarrow \mathcal{C}$ where p is a predicate in $P^{\leq k}$ and \mathcal{C} is the constraint satisfying p in M . Let P' be any set of clauses obtained from $P^{\leq k+1}$ by replacing some of the occurrences of predicates $p(X)$ from $P^{\leq k}$ in the bodies of clauses in $P^{\leq k+1}$ with their corresponding interpretation \mathcal{C} in M . Then*

1. *If $P^{\leq k+1}$ has a model then so does P' ;*
2. *If P' contains no predicate from $P^{\leq k}$, then P' is linear.*

4 Algorithm for solving sets of non-linear Horn clauses

A basic procedure for solving a set of non-linear Horn clauses using a linear Horn clause solver is presented in Algorithms 1 and 2. From now on we use linear solver for linear Horn clause solver for simplicity. The main procedure **SOLVE**(P) takes a set of non-linear Horn clauses P as input and outputs upon termination (*safe, solution*) if P is solvable or (*unsafe, counterexample*) otherwise. We represent counterexample as a *trace tree*. For a linear program it corresponds to a sequence of clauses.

Definition 4 (S_t) *Let S be an interpretation of a set of Horn clauses P . Let t be any trace tree of P and T_t be its corresponding Horn clause derivation. Then we denote by A_t the set of non-constraint atom label of T_t . Then S_t is defined to be the set*

$$S_t = \{(H \leftarrow \mathcal{C}) \mid (H \leftarrow \mathcal{C}) \in S \wedge H \notin A_t\}.$$

Algorithm 2 is an extended version of Algorithm 1, which uses the solution for $P^{\leq k}$ to help to linearise $P^{\leq k+1}$ and also allows a more refined termination condition based on whether or not the solution for $P^{\leq k}$ is used in constructing a counterexample for $P^{\leq k+1}$.

The procedures make use of several sub-procedures which will be described next.

4.1 Components of the algorithm

- **KDIM**(P, k): produces an at-most- k -dimension program $P^{\leq k}$ (Definition 2). By definition, $P^{\leq k}$ is linear for $k = 0$. For our example program presented in Figure 1, $Fib^{\leq 0}$ is shown in Figure 3 which is linear since there is at-most one non-constraint atom in the body of each clauses.

- **SOLVE_LINEAR(P')**: solves a set of linear Horn clauses P' . We assume the following about a linear solver: (i) if it terminates on P' , then it returns either *safe* and a *solution* or *unsafe* and a *counterexample*; (ii) it is sound, that is, if it returns a *solution* S for P' then P' has a model and S is a solution (model) of P' ; if it returns *unsafe* and a counterexample cEx then P' is unsafe and cEx is a witness. In our setting (Algorithms 1 and 2), P' corresponds to a linearised version of $P^{\leq k}$ for some P and $k \geq 0$. For technical reasons, the top level predicate $false[k]$ of $P^{\leq k}$ if any, is renamed to *false* before passing to a linear solver.

In essence, any Horn clause solver which complies with our assumption, for example QARMC [17], Convex polyhedral analyser [23], ELDARICA [20] etc. can be used in a black-box fashion but in this paper, we make use of a solver described in [23], which is based on abstract interpretation [6] over the domain of convex polyhedra [7] but without refinement using finite tree automata. The solver produces the following solution for the program in Figure 3. We can check it is in fact a solution (model).

```
fib(0)(A,B) :- [-A>= -1,A>=0,B=A].
fib[0](A,B) :- [-A>= -1,A>=0,B=A].
false[0] :- <>. % <> means that there is no model for false[0],
%so we can discard it
```

- **LINEARISE(P,k,S)** generates a linear set of clauses from $P^{\leq k}$ and an interpretation S for $P^{\leq k}$. Let S be a set of constrained facts of the form $p(X) \leftarrow \mathcal{C}$, where p is a predicate from $P^{\leq k}$, the procedure replaces every clause from $P^{\leq k}$ with head $p(X)$ by $p(X) \leftarrow \mathcal{C}$. This produces a set of clauses say P' . Then the procedure **LINEARISE_PE(P',M)** is called, which is the linearisation procedure based on partial evaluation described in Section 3 where M is a bound for the stack usage for linearising $P^{\leq k}$.

An excerpt from $Fib^{\leq 1}$ is shown below.

```
false(1) :- A>5, B<A, fib(1)(A,B).
fib(1)(A,B) :- A>1, C=A-2, E=A-1, B=F+D, fib(1)(C,D), fib[0](E,F).
fib(0)(A,B) :- B=A, A<1, A>=0.
```

After reusing the solution obtained for $Fib^{\leq 0}$ and linearising, we obtain the following set of linear clauses.

```
false(1) :- A>5, B<A, fib(1)(A,B).
fib(1)(A,B) :- -A>= -2, A>1, A-C=2, B-D=1, fib(1)(C,D).
```

Continuing to run our algorithm, the following solution obtained for $Fib^{\leq 2}$ becomes a solution for the program in Figure 1 (the original program) and the algorithm terminates.

```
fib(0)(A,B) :- [-A>= -1,A>=0,B=1].
fib[0](A,B) :- [-A>= -1,A>=0,B=1].
fib(1)(A,B) :- [A>=2,A+ -B=0].
fib[1](A,B) :- [A+ -B>= -1,B>=1,-A+B>=0].
fib(2)(A,B) :- [A>=4,-2*A+B>= -3].
fib[2](A,B) :- [A>=0,B>=1,-A+B>=0].
```


Algorithm 1: Algorithm for solving a set of Horn clauses

```

1 Procedure SOLVE( $P$ )
  Input: A set of CHCs  $P$ 
  Output: ( $safe, solution$ ), ( $unsafe, cex$ )
2  $k \leftarrow 0$ ;
3  $P' \leftarrow \text{LINEARISE}(P, k, \emptyset)$ ;
4 ( $status, Result$ )  $\leftarrow \text{SOLVE\_LINEAR}(P')$ ;      /* Result is a solution or a cex */
5 if  $status = safe$  then
6   | if ( $Result^{\uparrow P}$  is a solution of  $P$ ) then return ( $safe, Result^{\uparrow P}$ );
7   |  $k \leftarrow k + 1$ ;
8 else
9   | return ( $unsafe, Result$ );                      /* Result is a cex */
10 goto 3;

```

Algorithm 2: Algorithm for solving a set of Horn clauses, with reuse of lower dimension solutions

```

1 Procedure SOLVE( $P$ )
  Input: A set of CHCs  $P$ 
  Output: ( $safe, solution$ ), ( $unsafe, cex$ )
2  $k \leftarrow 0$ ;
3  $S \leftarrow \emptyset$ ;
4  $P' \leftarrow \text{LINEARISE}(P, k, S)$ ;
5 ( $status, Result$ )  $\leftarrow \text{SOLVE\_LINEAR}(P')$ ;      /* Result is a solution or a cex */
6 if  $status = safe$  then
7   | if ( $Result^{\uparrow P}$  is a solution of  $P$ ) then return ( $safe, Result^{\uparrow P}$ );
8   |  $k \leftarrow k + 1$ ;
9   |  $S \leftarrow Result$ ;
10 else
11   | if  $S \cap Result|_{cEx} = \emptyset$  then return ( $unsafe, Result$ ); /* Result is a linear cex */
12   |
13   |  $S \leftarrow S \setminus Result|_{cEx}$ ;          /*  $Result|_{cEx}$ : Definition 4 */
14 goto 4;

```

Algorithm 3: Algorithm for linearising a set of clauses

```

1 Procedure LINEARISE( $P, k, S$ )
  Input: A set of CHCs  $P$ , an integer  $k$  and a set of constrained facts  $S$ 
  Output: A linearised set of clauses  $P_{lin}$ 
2  $P^{\leq k} \leftarrow \text{KDIM}(P, k)$ ;                      /* Definition 2 */
3  $P' \leftarrow \text{SUBSTITUTE}(P^{\leq k}, S)$ ;          /* substitute atoms of  $P^{\leq k}$  with their
   interpretations from  $S$  */
4  $M \leftarrow (i - 1) * k + 1$ ;      /* where  $i$  is the maximal number of body atoms of  $P$  */
5  $P_{lin} \leftarrow \text{LINEARISE\_PE}(P', M)$ ;          /* Section 3.1 */
6 return  $P_{lin}$ 

```


Proposition 3 (Soundness) *If Algorithm 1 or 2 returns safe and a solution S for a set of clauses P then P is safe and S is in fact a solution of P ; if it returns unsafe and a counterexample cEx then P is unsafe and cEx is a witness.*

Another desired property of the Algorithm 2 is that of progress, that is, the same counterexample does not arise more than once.

5 Experimental results

We made a prototype implementation of Algorithm 2 in the tool called *LHornSolver*¹. It uses the solver described in [23] without refinement as a linear Horn clause solver. *LHornSolver* is written in Ciao prolog [19] and is interfaced with Parma Polyhedra Library [2] and Yices SMT solver [10] for handling constraints. Then we carried out an experiment on a set of 44 non-linear CHC verification problems taken from the repository² of software verification benchmarks, the recursive category of SV-COMP³ [4] and the tool QARMC. The experiments were run on a MAC computer running OS X on 2.3 GHz Intel core i7 processor and 8 GB memory. The benchmarks that we use in the experiments are not beyond the capabilities of existing solvers, but they are challenging. These programs are first translated to Prolog syntax using the tools Eldarica⁴ [20] and SeaHorn [18]. Our aim with these experiments is to explore: (1) whether using a linear solver for non-linear problem solving is practical; (2) the relationship between the solvability of a problem and its dimension; and (3) how the current results compare with the results using the state of the art non-linear Horn clause verification tool (in our case RAHFT [23]). The results are summarized in Table 1.

In the table *Program* represents a program, *Safety* represents a verification result, *# Itr.* and *Time (s)* successively represent the number of refinement iterations and the time in seconds need to solve a program using both RAHFT and LHornSolver. It is to note that the underlying abstract interpreter, that is, the convex polyhedral analyser (CPA) is the same for both *RAHFT* and *LHornSolver* but *LHornSolver* uses it to solve linear Horn clauses though the CPA is not optimised for linear problems. The column *# Itr.* for LHornSolver represents a value of k for which a solution of $P^{\leq k}$ (under-approximation) of a set of clauses P becomes a solution for P or $P^{\leq k}$ becomes unsafe. The symbol “?” means that the result is unknown within the given time bound. The result “safe” means that the program is safe (solvable) and “unsafe” means it is unsafe.

LHornSolver solves 27 out of 44 (about 61%) problems within a second. In most of these problems, a solution of an under approximation ($P^{\leq k}$) becomes a solution for the original program or $P^{\leq k}$ becomes unsafe for a fairly small value of k (1 or 2). This suggests that the solvability of a problem is shallow with respect to its dimension. This demonstrates the feasibility of solving a set of non-linear Horn clauses using a solver for linear Horn clauses.

In contrast, RAHFT solves all the problem. The difference in results maybe due to the following reason: the linear solver that is used in LHornSolver is the CPA (without refinement in contrast to [23]). The solver terminates but produces *false alarms*. If we use CPA with refinement as in [23], then we lose predicates names (due to program transformation), so the solution or counterexamples produced by the tool do not correspond to the original program (it is very hard to keep track of the changes). This hinders the reuse of solution from lower dimension to linearise program of higher dimension or refine it

¹<https://github.com/bishoksan/LHornSolver>

²<https://svn.sosy-lab.org/software/sv-benchmarks/trunk/clauses/LIA/Eldarica/RECUR/>

³<http://sv-comp.sosy-lab.org/2015/benchmarks.php>

⁴<https://github.com/uuverifiers/eldarica>

	RAHFT			LHornSolver		
Program	Safety	# iter.	Time (s)	Safety	#iter.	Time (s)
Addition03_false-unreach	safe	2	< 1	?	?	?
McCarthy91_false-unreach	unsafe	0	< 1	?	?	?
addition.nts.pl	safe	0	< 1	safe	1	< 1
bfprt.nts.pl	safe	0	< 1	safe	2	4
binarysearch.nts.pl	safe	0	< 1	safe	1	1.1
countZero.nts.pl	safe	0	< 1	safe	1	< 1
eq.horn	unsafe	0	< 1	unsafe	2	< 1
fib.pl	safe	0	< 1	?	?	?
identity.nts.pl	safe	0	< 1	safe	1	< 1
merge.nts.pl	safe	0	< 1	safe	1	1.7
palindrome.nts.pl	safe	0	< 1	safe	1	< 1
parity.nts.pl	unsafe	1	< 1	?	?	?
remainder.nts.pl	unsafe	0	< 1	unsafe	1	< 1
revlen.pl	safe	0	< 1	safe	1	< 1
running.nts.pl	unsafe	1	< 1	?	?	?
sum_10x0_false-unreach	unsafe	10	10	?	?	?
sum_non_eq_false-unreach	unsafe	0	< 1	?	?	?
sumal.horn	unsafe	0	< 1	unsafe	1	< 1
suma2.horn	unsafe	0	< 1	unsafe	2	< 1
summ_SG1.r.horn	safe	0	< 1	?	?	?
summ_SG2.r.horn	safe	8	78	?	?	?
summ_SG3.horn	safe	0	< 1	safe	1	< 1
summ_b.horn	safe	2	1.7	?	?	?
summ_binsearch.horn	safe	1	3	?	?	?
summ_cil.casts.horn	safe	0	< 1	safe	1	< 1
summ_formals.horn	safe	0	< 1	safe	1	< 1
summ_g.horn	safe	0	< 1	?	?	?
summ_globals.horn	safe	0	< 1	safe	1	< 1
summ_h.horn	safe	0	< 1	safe	2	< 1
summ_local-ctx-call.horn	safe	0	< 1	safe	1	< 1
summ_locals.horn	safe	0	< 1	?	?	?
summ_locals2.horn	safe	0	< 1	safe	1	< 1
summ_locals3.horn	safe	0	< 1	safe	1	< 1
summ_locals4.horn	safe	0	< 1	safe	2	2.2
summ_mccarthy2.horn	safe	3	5	?	?	?
summ_multi-call.horn	safe	0	< 1	safe	1	< 1
summ_nested.horn	safe	0	< 1	safe	1	< 1
summ_ptr_assign.horn	safe	0	< 1	safe	1	< 1
summ_recursive.horn	safe	0	< 1	?	?	?
summ_rholocal.horn	safe	0	< 1	safe	1	< 1
summ_rholocal2.horn	safe	0	< 1	safe	1	< 1
summ_slicing.horn	safe	0	< 1	?	?	?
summ_sums.horn	safe	0	< 1	?	?	?
summ_typedef.horn	safe	0	< 1	safe	1	< 1
summ_x.horn	safe	0	< 1	?	?	?
Average		0.64	2.3		1.185	< 1

Table 1: Experimental results on non-linear CHC verification problems with a timeout of 5 minutes.

using the counterexample trace. Other solvers which don't modify the programs but produce solutions or counterexamples can be used as a linear solver in principle and we leave it for the future work. Another disadvantage of using CPA is that, if it cannot solve a linear program, then it emits an abstract trace which is checked for a feasibility. If it is spurious then `LHornSolver` returns with *unknown* (in principle we can refine the program but the refinement will have the problem as mentioned above). So it is highly unlikely that the trace picked by the tool non-deterministically results to be a real counterexample. We noticed in our experiments that the trace picked was spurious most of the times and *LHornSolver* immediately returned “unknown” answer before the timeout. This also explains why solving time of *LHornSolver* is less than that of *RAHFT*.

The interpreter described in Figure 6 computed a *permutation* of the atoms in a clause body; partial evaluation of the permutation procedure can cause a blow-up of the size of the linearised program, relative to the number of atoms in clause bodies. In our experiment because the maximum number of atoms in the bodies of the clauses in our benchmark programs was 5 and the value of k was relatively small ($k = 0 \dots 2$). The permutation procedure can be avoided if we first generate an at-most- k -dimension program whose body atoms are ordered by increasing dimension. This needs unfolding of the ε -clauses, since atoms whose predicate is $p^{\leq d}$ cannot be ordered directly; only atoms with predicates of the form $p^=d$ can be ordered. We have not yet evaluated the trade-offs in these two approaches.

6 Related Work

In the world of Horn clause solvers, after fixing a constraint theory, we can distinguish solvers depending on whether they can handle general non-linear Horn clauses or not. A majority of solvers [18, 16, 30, 29, 23] handle non-linear Horn clauses but there are notable exceptions like VeriMAP [8] or Sally⁵. For both VeriMAP and Sally, their underlying reasoning engine handles only linear Horn clauses which restricts, in principle, their applicability. Our contribution is to lift this restriction by allowing those tools to be applied on arbitrary sets of Horn Clauses, linear or not, through a linearisation procedure that underapproximates the set of solutions. We give empirical evidence that this underapproximation often provides enough coverage to enable the verification of the original set of Horn clauses. To summarize, we allow solvers with restrictions to be applied on any input at the price of an under-approximation which often results in full coverage.

Our linearisation method based on partial evaluation described in Section 3.1 is related to the linearisation method based on *fold-unfold transformations* described by De Angelis *et al.* [9]. While their procedure transforms the target set of clauses directly, we transform an interpreter for the clauses using a generic partial evaluation procedure. Any clause transformation procedure could be formulated as a meta-program and partial evaluation applied to that program to yield the specified transformation. Thus neither approach offers any more power than the other. However the use of partial evaluation is arguably more flexible. The interpreter that is partially evaluated in our procedure is a standard interpreter for Horn clauses, modified with a bound on the size of goals, directly incorporating a general result that there is an upper bound on the size of goals in derivations with dimension-bounded programs. This provides a very generic starting point for the transformation with an explicit relation to the semantics of the clauses. A whole family of similar transformations could be formulated by varying the interpreter (for example using breadth-first search). The procedure in [9] is tailored to a restrictive setting where only goal clauses (integrity constraints) are non-linear and rest of the clauses are linear; correctness has to be established for that case.

⁵<https://github.com/SRI-CSL/sally>

Ganty, Iosif and Konečný [15] used the notion of *tree dimension* for computing summaries of procedural programs by underapproximating them. Roughly speaking, they compute procedure summaries iteratively, starting from the program behaviors captured by derivation trees of dimension 0. Then they reuse these summaries to compute summaries for program behaviors captured by derivation trees of dimension 1 and so on for 2, 3, etc. Kafle, Gallagher and Ganty [24] adapted the idea of dimension-based underapproximations to the setting of Horn clause systems. They gave empirical evidence supporting the thesis that for small values of the dimension the solutions are general enough to hold for every dimension. Their approach still required the use of general Horn-clause solvers capable of handling non-linear clauses. In this paper, we lift this requirement and allow the use of solvers for linear clauses only. Moreover, we provide an abstraction refinement loop that enables the solutions for lower dimension to be reused when searching for solutions in higher dimension.

7 Conclusion and future work

We presented an *abstraction-refinement* approach for solving a set of non-linear Horn clauses using an off-the-shelf linear Horn clause solver. It was achieved through a linearisation of a dimension bounded set of Horn clauses (which are known to be linearisable) using partial evaluation and the use of a linear Horn clause solver. Experiment on a set of non-linear Horn clause verification problems using our approach shows that the approach is feasible (a linear solver can be used for solving non-linear problems) and the solvability of a problem is shallow with respect to its dimension.

In the future, we plan to compare our results with the results from a specialised linear Horn clause solver like VeriMap and other non-linear Horn clause solvers. We also plan to experiment with different linearisation strategies for Horn clauses and study their effects in Horn clause verification.

Acknowledgement

The authors would like to thank José F. Morales for his help with Ciao Prolog foreign language interface and some parts of the implementation.

References

- [1] F. N. Afrati, M. Gergatsoulis, and F. Toni. Linearisability on datalog programs. *Theor. Comput. Sci.*, 308(1-3):199–226, 2003.
- [2] R. Bagnara, P. M. Hill, and E. Zaffanella. 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, 2008.
- [3] C. Baier and C. Tinelli, editors. *Tools and Algorithms for the Construction and Analysis of Systems - 21st International Conference, TACAS 2015, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2015, London, UK, April 11-18, 2015. Proceedings*, volume 9035 of *Lecture Notes in Computer Science*. Springer, 2015.
- [4] D. Beyer. Software verification and verifiable witnesses - (report on SV-COMP 2015). In Baier and Tinelli [3], pages 401–416.
- [5] N. Bjørner, K. L. McMillan, and A. Rybalchenko. On solving universally quantified Horn clauses. In F. Logozzo and M. Fähndrich, editors, *SAS*, volume 7935 of *LNCS*, pages 105–125. Springer, 2013.

- [6] P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In R. M. Graham, M. A. Harrison, and R. Sethi, editors, *POPL*, pages 238–252. ACM, 1977.
- [7] P. Cousot and N. Halbwachs. Automatic discovery of linear restraints among variables of a program. In *Proceedings of the 5th Annual ACM Symposium on Principles of Programming Languages*, pages 84–96, 1978.
- [8] E. De Angelis, F. Fioravanti, A. Pettorossi, and M. Proietti. Verimap: A tool for verifying programs through transformations. In E. Ábrahám and K. Havelund, editors, *TACAS*, volume 8413 of *LNCS*, pages 568–574. Springer, 2014.
- [9] E. De Angelis, F. Fioravanti, A. Pettorossi, and M. Proietti. Proving correctness of imperative programs by linearizing constrained Horn clauses. *TPLP*, 15(4-5):635–650, 2015.
- [10] B. Dutertre. Yices 2.2. In A. Biere and R. Bloem, editors, *Computer-Aided Verification (CAV'2014)*, volume 8559 of *Lecture Notes in Computer Science*, pages 737–744. Springer, July 2014.
- [11] J. Esparza, P. Ganty, S. Kiefer, and M. Luttenberger. Parikh’s theorem: A simple and direct automaton construction. *Inf. Process. Lett.*, 111(12):614–619, 2011.
- [12] J. Esparza, S. Kiefer, and M. Luttenberger. On fixed point equations over commutative semirings. In *STACS 2007, 24th Annual Symposium on Theoretical Aspects of Computer Science, Proceedings*, volume 4393 of *LNCS*, pages 296–307. Springer, 2007.
- [13] J. P. Gallagher. Transforming logic programs by specialising interpreters. In *Proceedings of the 7th European Conference on Artificial Intelligence (ECAI-86), Brighton*, pages 109–122, 1986.
- [14] J. P. Gallagher. Tutorial on specialisation of logic programs. In D. A. Schmidt, editor, *Proceedings of the ACM SIGPLAN Symposium on Partial Evaluation and Semantics-Based Program Manipulation, PEPM’93, Copenhagen, Denmark, June 14-16, 1993*, pages 88–98. ACM, 1993.
- [15] P. Ganty, R. Iosif, and F. Konečný. Underapproximation of procedure summaries for integer programs. In N. Piterman and S. A. Smolka, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 19th International Conference, TACAS 2013, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2013, Rome, Italy, March 16-24, 2013. Proceedings*, volume 7795 of *Lecture Notes in Computer Science*, pages 245–259. Springer, 2013.
- [16] S. Grebenschikov, A. Gupta, N. P. Lopes, C. Popeea, and A. Rybalchenko. HSF(C): A software verifier based on Horn clauses - (competition contribution). In C. Flanagan and B. König, editors, *TACAS*, volume 7214 of *LNCS*, pages 549–551. Springer, 2012.
- [17] S. Grebenschikov, N. P. Lopes, C. Popeea, and A. Rybalchenko. Synthesizing software verifiers from proof rules. In J. Vitek, H. Lin, and F. Tip, editors, *ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI ’12, Beijing, China - June 11 - 16, 2012*, pages 405–416. ACM, 2012.
- [18] A. Gurfinkel, T. Kahsai, and J. A. Navas. SeaHorn: A framework for verifying C programs (competition contribution). In Baier and Tinelli [3], pages 447–450.
- [19] M. V. Hermenegildo, F. Bueno, M. Carro, P. López-García, E. Mera, J. F. Morales, and G. Puebla. An overview of ciao and its design philosophy. *TPLP*, 12(1-2):219–252, 2012.
- [20] H. Hojjat, F. Konečný, F. Garnier, R. Iosif, V. Kuncak, and P. Rümmer. A verification toolkit for numerical transition systems - tool paper. In D. Giannakopoulou and D. Méry, editors, *FM 2012: Formal Methods - 18th International Symposium, Paris, France, August 27-31, 2012. Proceedings*, volume 7436 of *Lecture Notes in Computer Science*, pages 247–251. Springer, 2012.
- [21] N. Jones, C. Gomard, and P. Sestoft. *Partial Evaluation and Automatic Software Generation*. Prentice Hall, 1993.
- [22] N. D. Jones. Transformation by interpreter specialisation. *Sci. Comput. Program.*, 52:307–339, 2004.
- [23] B. Kafle and J. P. Gallagher. Horn clause verification with convex polyhedral abstraction and tree automata-based refinement. *Computer Languages, Systems & Structures*, 2015.

- [24] B. Kafle, J. P. Gallagher, and P. Ganty. Decomposition by tree dimension in Horn clause verification. In A. Lisitsa, A. P. Nemytykh, and A. Pettorossi, editors, *Proceedings of the Third International Workshop on Verification and Program Transformation, VPT 2015, London, United Kingdom, 11th April 2015.*, volume 199 of *EPTCS*, pages 1–14, 2015.
- [25] M. Leuschel. Partial evaluation of the "real thing". In L. Fribourg and F. Turini, editors, *Logic Programming Synthesis and Transformation - Meta-Programming in Logic. 4th International Workshops, LOPSTR'94 and META'94, Pisa, Italy, June 20-21, 1994, Proceedings*, volume 883 of *Lecture Notes in Computer Science*, pages 122–137. Springer, 1994.
- [26] M. Leuschel, D. Elphick, M. Varea, S.-J. Craig, and M. Fontaine. The Ecce and Logen partial evaluators and their web interfaces. In J. Hatcliff and F. Tip, editors, *PEPM 2006*, pages 88–94. ACM, 2006.
- [27] M. Leuschel and G. Vidal. Fast offline partial evaluation of logic programs. *Inf. Comput.*, 235:70–97, 2014.
- [28] M. Luttenberger and M. Schlund. Convergence of newton's method over commutative semirings. *Inf. Comput.*, 246:43–61, 2016.
- [29] K. L. McMillan and A. Rybalchenko. Solving constrained Horn clauses using interpolation. Technical report, Microsoft Research, 2013.
- [30] P. Rümmer, H. Hojjat, and V. Kuncak. Disjunctive interpolants for Horn-clause verification. In N. Sharygina and H. Veith, editors, *CAV*, volume 8044 of *Lecture Notes in Computer Science*, pages 347–363. Springer, 2013.