

Solving Recursion-Free Horn Clauses over LI+UIF

Ashutosh Gupta^{1,2}, Corneliu Popeea², and Andrey Rybalchenko²

¹IST Austria ²Technische Universität München

Abstract. Verification of programs with procedures, multi-threaded programs, and higher-order functional programs can be effectively automated using abstraction and refinement schemes that rely on spurious counterexamples for abstraction discovery. The analysis of counterexamples can be automated by a series of interpolation queries, or, alternatively, as a constraint solving query expressed by a set of recursion free Horn clauses. (A set of interpolation queries can be formulated as a single constraint over Horn clauses with linear dependency structure between the unknown relations.) In this paper we present an algorithm for solving recursion free Horn clauses over a combined theory of linear real/rational arithmetic and uninterpreted functions. Our algorithm performs resolution to deal with the clausal structure and relies on partial solutions to deal with (non-local) instances of functionality axioms.

1 Introduction

Constraint solving is a vehicle of software verification that provides symbolic reasoning techniques for dealing with assertions describing program behaviors. In particular, abstraction and refinement techniques greatly benefit from applying constraint solving, where interpolation techniques [1–3, 5, 6, 11–13, 16–19] play a prominent role today. Roughly, interpolation computes an assertion that separates two mutually unsatisfiable assertions and only refers to their shared symbols.

Certain abstraction refinement tasks cannot be *directly* expressed as an interpolation question. For example, abstraction refinement for imperative programs with procedures [12] and for higher order functional programs [14, 20], require additional pre-processing that splits discovered spurious counterexamples in multiple ways and applies interpolation on each splitting. Alternatively, as exemplified by an abstraction refinement procedure for multi-threaded programs [9], this preprocessing and series of interpolation computations can be expressed using a single constraint that consists of a finite set of recursion-free Horn clauses interpreted over the logical theory that is used to describe program behaviors.

In this paper, we present an algorithm for solving Horn clauses over a combination of linear rational/real arithmetic, uninterpreted functions and queries. Our algorithm opens new possibilities for the development of abstraction refinement schemes by providing the verification method designer an expressive,

declarative way to specify what the refinement procedure needs to compute using Horn clauses. Several existing abstraction refinement schemes can directly benefit from our algorithm, e.g., for programs with procedures [11,12], for multi-threaded programs [9], and for higher-order functional programs [14,20,21].

Related work In [9] we presented an algorithm that deals with recursion-free Horn clauses over linear real/rational arithmetic. Here, we present an extension with uninterpreted functions.

Technically, our treatment of uninterpreted functions can be seen as a generalization of partial interpolants [17] to partial solutions for recursion-free Horn clauses, i.e., clauses that do not have cyclic dependencies between the occurring queries. Our algorithm follows a general scheme of combining interpolation procedures for different theories [8,22].

The following example illustrates the relation to interpolation. First, we consider an interpolation question for a pair of mutually unsatisfiable assertions $a(x, y)$ and $b(y, z)$ in a logical theory. An interpolant is an assertion $I(y)$ such that $I(y)$ is a logical consequence of $a(x, y)$, $I(y)$ and $b(y, z)$ are mutually unsatisfiable, and $I(y)$ only contains non-theory constants that are shared by $a(x, y)$ and $b(y, z)$, which is y in our example. Now, we present our re-formulation of interpolation as a constraint solving question for constraints given by recursion-free Horn clauses. We introduce a relation $Q_I(y)$ that represents an interpolant that we want to compute. We represent the interpolation conditions by the following two Horn clauses: $a(x, y) \rightarrow Q_I(y)$ and $Q_I(y) \wedge b(y, z) \rightarrow false$. Any interpretation of $Q_I(y)$ that only refers to y (and theory constants) is an interpolant for $a(x, y)$ and $b(y, z)$. In Section 6, we discuss the relation of this paper with [17] in more detail.

Horn clauses with more than two unknown queries in the body do not directly correspond to interpolation problems. For example, we consider two relations $Q_I(y)$ and $Q_J(y)$ that represent assertions we want to compute together with the Horn clauses $a(x, y) \rightarrow Q_I(y)$, $b(y, z) \rightarrow Q_J(y)$, and $Q_I(y) \wedge Q_J(y) \rightarrow false$. Solving these clauses using interpolation requires two invocations of an interpolation procedure (i.e., interpolation between $a(x, y)$ and $b(y, z)$ determines $Q_I(y)$ interpolation between $b(y, z)$ and $Q_I(y)$ determines $Q_J(y)$) that we would like to avoid for efficiency considerations. Furthermore, by computing $Q_I(y)$ and $Q_J(y)$ one after the other it is not evident how to compute solutions satisfying certain preference conditions, e.g., where all constraints are within predefined bounds (such conditions are useful for abstraction refinement, as shown by the FOCI procedure [15]).

Organization Section 2 illustrates our algorithm. Section 3 provides formal definitions. We present the solving algorithm in Section 4 and discuss its correctness and complexity in Section 5. Section 6 concludes and clarifies the connection of our algorithm to interpolation procedures.

We also provide an extended version of this paper [10]. In this extended version, Appendix A illustrates how Horn clauses can be used for abstraction refinement of procedural and multi-threaded programs, Appendix B contains a

complete example execution of our algorithm, and Appendix C presents proofs of the key theorems.

2 Illustration

In this section, we shall illustrate our proposed algorithm by solving an example set of Horn clauses. The example set of Horn clauses \mathcal{HC} is presented in Figure 1(a) and consists of three clauses. Our algorithm is looking for solutions to the two symbols $S(\mathbf{t}, \mathbf{u}, \mathbf{v})$ and $E(\mathbf{t}, \mathbf{u})$ that we name *queries*. To obtain solutions over the domain of linear arithmetic and uninterpreted functions, our algorithm proceeds following three steps.

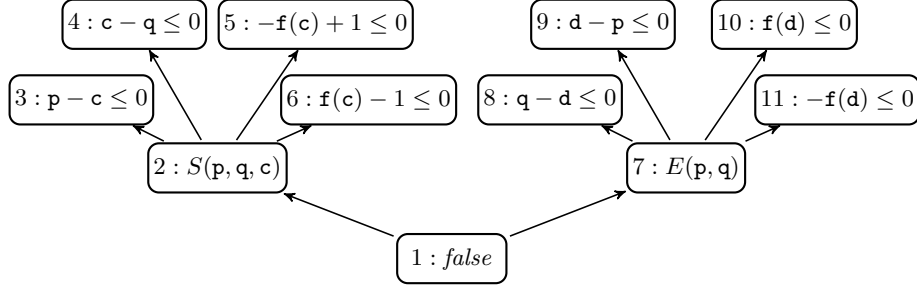
Resolution tree Our solving algorithm starts by constructing from \mathcal{HC} a resolution tree R shown in Figure 1(b). We label nodes of R with indices for easy reference. From the first clause, the algorithm constructs the subtree rooted at label 2. In this subtree, we have edges between the node corresponding to the head of the clause (labeled 2) and the nodes corresponding to the body of the clause (labeled 3–6). A second subtree rooted at the node labelled 7 is constructed from the second clause. With the appearance of the queries $S(\mathbf{t}, \mathbf{u}, \mathbf{v})$ and $E(\mathbf{t}, \mathbf{u})$ in the body of the third clause from \mathcal{HC} , the two previously constructed subtrees are extended in a tree with the root labeled corresponding to the clause head, (1 : *false*). The extension of the subtrees leads to the variables occurring in these subtrees to be renamed to a common set of variables $\mathbf{p}, \mathbf{q}, \mathbf{c}$. Note that, the set of clauses \mathcal{HC} is satisfiable, and, consequently, the conjunction of the predicates from the leaves of the resolution tree is unsatisfiable.

Proof tree Next, our algorithm constructs a proof tree that proves unsatisfiability of the constraints from the leaves of the resolution tree. For the resolution tree from Figure 1(b), our algorithm computes the proof tree P shown in Figure 1(c). A linear combination rule is applied to derive the constraint $(\mathbf{c} - \mathbf{d} \leq 0)$ from the premises $(\mathbf{c} - \mathbf{q} \leq 0)$ and $(\mathbf{q} - \mathbf{d} \leq 0)$. The linear combination rule is also used to derive $(\mathbf{d} - \mathbf{c} \leq 0)$ from the premises $(\mathbf{p} - \mathbf{c} \leq 0)$ and $(\mathbf{d} - \mathbf{p} \leq 0)$. A congruence rule is used to relate function symbols applied to equivalent arguments. This rule derives $(\mathbf{f}(\mathbf{c}) - \mathbf{f}(\mathbf{d}) \leq 0)$ from the premises $(\mathbf{c} - \mathbf{d} \leq 0)$ and $(\mathbf{d} - \mathbf{c} \leq 0)$. Lastly, $(1 \leq 0)$ is derived by applying the linear combination rule on three premises, $(\mathbf{f}(\mathbf{d}) \leq 0)$, $(\mathbf{f}(\mathbf{c}) - \mathbf{f}(\mathbf{d}) \leq 0)$, and $(-\mathbf{f}(\mathbf{c}) + 1 \leq 0)$.

Partial and final solutions The proof tree P explicates the inference rules and the order in which to apply them to derive the false constraint $(1 \leq 0)$. The main idea behind our solving algorithm is to apply corresponding inference rules in the same order to derive a solution for the Horn clauses. We obtain an annotated proof tree (see Figure 1(d)) where for each of the premises used in P , our algorithm creates one tree with the same number of nodes as R . We call these trees, which are annotated with formulas that will be explained next, partial-solution trees.

$$\mathcal{HC} = \{ \forall p, q, c : p \leq c \wedge c \leq q \wedge -f(c) + 1 \leq 0 \wedge f(c) - 1 \leq 0 \rightarrow S(p, q, c), \\ \forall r, s, d : s \leq d \wedge d \leq r \wedge f(d) \leq 0 \wedge -f(d) \leq 0 \rightarrow E(r, s), \\ \forall t, u, v : S(t, u, v) \wedge E(t, u) \rightarrow \text{false} \}$$

(a)



(b)

$$\frac{\frac{\frac{c - q \leq 0 \quad q - d \leq 0}{c - d \leq 0} \quad \frac{p - c \leq 0 \quad d - p \leq 0}{d - c \leq 0}}{f(d) \leq 0 \quad f(c) - f(d) \leq 0} \quad -f(c) + 1 \leq 0}{1 \leq 0}$$

(c)

$$\frac{\frac{\frac{c - q \leq 0[I_1] \quad q - d \leq 0[I_2]}{c - d \leq 0[I_3]} \quad \dots}{f(d) \leq 0[\dots] \quad f(c) - f(d) \leq 0[\dots]} \quad -f(c) + 1 \leq 0[\dots]}{1 \leq 0[II]}$$

(d)

Fig. 1. (a) A set of Horn clauses \mathcal{HC} . (b) Corresponding resolution tree R . (c) Proof of unsatisfiability P for the constraints from the leaves of the resolution tree. For abbreviation, we did not mark nodes of subtree of $f(c) - f(d) \leq 0$ with the applied proof rules. (d) A part of the annotated proof tree. The partial solutions I_1 , I_2 , I_3 , and II are presented in Figure 2.

The tree I_1 corresponds to the premise $(c - q \leq 0)$, I_2 corresponds to the premise $(q - d \leq 0)$ and both trees are shown in Figure 2. Two or more premises are used to derive a new fact in the proof tree and, likewise, two or more corresponding partial-solution trees are used to derive a new tree using a specific inference rule. The two trees I_1 and I_2 shown in the top part of Figure 2




Fig. 2. Four partial-solution trees Π_1 , Π_2 , Π_3 , and Π . Π_1 and Π_2 are derived from the nodes $(c - q \leq 0)$ and $(q - d \leq 0)$ from the proof tree P from Figure 1(d). Π_3 is obtained by applying a combination rule to Π_1 and Π_2 . Π annotates the false constraint $(1 \leq 0)$ from P and the final solution of \mathcal{HC} can be derived from Π . In particular, the nodes labeled “2” and “7” contain the solutions for $S(p, q, c)$ and $E(p, q)$, respectively.

are combined using a rule corresponding to the arithmetic combination rule. The rule takes a pair of corresponding nodes, one from Π_1 and one from Π_2 , and computes a node in the resulting tree Π_3 . For the node labeled $(2 : c - q \leq 0)$ from Π_1 and the node labeled $(2 : 0 \leq 0)$ from Π_2 , the algorithm adds the two constraints and creates a node labeled $(2 : c - q \leq 0)$ in Π_3 . Similarly, the nodes labeled $(1 : c - q \leq 0)$ and $(1 : q - d \leq 0)$ are used to obtain a node labeled $(1 : c - d \leq 0)$ in Π_3 .

Following the derivation of the proof tree P , inference rules are used to combine partial-solution trees until a final-solution tree corresponding to the rule applied at the bottom of the proof tree. The final-solution tree is Π and is shown in Figure 2. The node labeled “2” contains the solution for $S(p, q, c)$ and it can be simplified to $S(p, q, c) = (p < q \vee p \leq q \wedge f(p) \geq 1)$. The solution from the node labeled “7” can be simplified to $E(p, q) = (p > q \vee p \geq q \wedge f(p) \leq 0)$. The solutions obtained for $S(p, q, c)$ and $E(p, q)$ indeed satisfy the set of Horn clauses \mathcal{HC} from Figure 1(a).

3 Recursion-free Horn clauses

This section presents auxiliary definitions together with the syntax and semantics of recursion-free Horn clauses over linear arithmetic, uninterpreted functions, and queries.

Syntax We assume countable sets of *variables* \mathcal{V} , with $v \in \mathcal{V}$, *function symbols* \mathcal{F} , with $f \in \mathcal{F}$, and *predicate symbols* \mathcal{P} , with $p \in \mathcal{P}$. Let the arity of function and predicate symbols be encoded in their names. In addition, we assume a set of *number symbols* \mathcal{N} , with $\{0, n\} \subseteq \mathcal{N}$, and an inequality symbol \leq . Then, we define:

$$\begin{array}{ll}
\text{terms} \ni t ::= n \mid nv \mid t + t \mid f(t, \dots, t) & \text{bodies} \ni b ::= a \mid q \mid b \wedge b \\
\text{atoms} \ni a ::= t \leq 0 & \text{heads} \ni h ::= a \mid q \mid \text{false} \\
\text{queries} \ni q ::= p(v, \dots, v) & \text{Horn clauses} \ni s ::= b \rightarrow h
\end{array}$$

Without loss of generality, as justified later, we assume that all variables that occur in a query are distinct.

A set of Horn clauses defines a binary *dependency* relation on predicate symbols. A predicate symbol $p \in \mathcal{P}$ *depends* on a predicate symbol $p_i \in \mathcal{P}$ if there is a Horn clause $\dots \wedge p_i(\dots) \wedge \dots \rightarrow p(\dots)$, i.e., when p appears in the head of a clause that contains p_i in its body. A set of Horn clauses is *recursion-free* if the corresponding dependency relation does not contain any cycles. A set of Horn clauses is *tree-like* if the corresponding dependency relation defines a tree-like graph, i.e., when 1) each predicate symbol appears at most once in the set of bodies and at most once in the set of heads of the given clauses, 2) there is no clause with an atom in its head, 3) there is one clause whose head is *false*. For example, the set of clauses $\{p(v_1) \wedge p(v_2) \rightarrow q(v_1, v_2), q(v_3, v_4) \rightarrow \text{false}\}$ is not

tree-like since the predicate symbol p appears more than once in the body of the first clause.

For the rest of the presentation, we consider a finite set of Horn clauses \mathcal{HC} that satisfies the following conditions. First, we assume that each variable occurs in at most one clause and that all variables occurring in a query are distinct. These assumptions simplify our presentation and can be established by an appropriate variable renaming and additional (in)equality constraints. Furthermore, we assume that \mathcal{HC} is recursion-free and tree-like. The recursion-free assumption is critical for ensuring termination of the solving algorithm presented in this paper. The tree-like assumption simplifies our presentation without imposing any restrictions on the algorithm's applicability. Any finite set of recursion-free clauses can be transformed into the tree-like form. The solution for the computed tree-like form can be translated into the solution for the original set of clauses.

Finally, we define *constraints* together with a *conjunctive constraint* fragment below.

constraints $\ni c ::= a \mid \neg c \mid c \wedge c \mid c \vee c$ conjunctive constraints $\ni \hat{c} ::= a \mid \hat{c} \wedge \hat{c}$

Auxiliary definitions We assume the following standard functions. For dealing with trees, let $nodes(T)$ be the nodes of a tree T , $root(T)$ be the root node of T , $leaves(T)$ be the leaves of T , and $subtree(o, T)$ be the subtree of T rooted in its node o . Furthermore, let $subterms(C)$ be the subterms occurring in a constraint C and $atoms(C)$ be the atoms occurring in C . Let $sym(t)$ be the variables and uninterpreted function symbols occurring in a term t .

Let $match(p(v_1, \dots, v_n), p'(v'_1, \dots, v'_m))$ return a substitution $\{v_1 \mapsto v'_1, \dots, v_n \mapsto v'_m\}$ if $p = p'$ (and hence $n = m$). Thus, if a substitution σ is the result of $match(p'(v_1, \dots, v_n), p'(v'_1, \dots, v'_m))$ then $p'(v_1, \dots, v_n)\sigma = p'(v'_1, \dots, v'_m)$, i.e., by applying the substitution we equate the queries. For example, $match(p_1(v_1), q(v_2, v_3))$ is not defined, and $match(q(v_1, v_2), q(v_3, v_4)) = \{v_1 \mapsto v_3, v_2 \mapsto v_4\}$. We assume a canonical extension of the unifier application to constraints and their combination into sequences and sets.

Given two substitutions $\sigma_1 = \{v_1 \mapsto v'_1, \dots, v_n \mapsto v'_n\}$ and $\sigma_2 = \{w_1 \mapsto w'_1, \dots, w_m \mapsto w'_m\}$ over disjoint domains, i.e., $\{v_1, \dots, v_n\} \cap \{w_1, \dots, w_m\} = \emptyset$, we define a *combined* substitution $\sigma_1 + \sigma_2 = \{v_1 \mapsto v'_1, \dots, v_n \mapsto v'_n, w_1 \mapsto w'_1, \dots, w_m \mapsto w'_m\}$.

Semantics Let \models be the (logical) satisfaction relation for our constraints in the combined theory of linear real/rational arithmetic and uninterpreted functions. We write $\models c$ when c is a valid constraint.

Let Σ be a function from queries to constraints. We assume that in the domain of Σ no two queries have an equal predicate symbol, all queries have disjoint variables, and each query is mapped to a constraint whose free variables occur in the query. For example, consider $\Sigma = \{p(v_1) \mapsto (v_1 \geq 0), q(v_2, v_3) \mapsto (v_2 \leq f(v_3))\}$.

We use Σ function to transform the set of Horn clauses containing queries into a set of query-free clauses as follows. In each clause $s \in \mathcal{HC}$ we replace

```

algorithm SOLVEHORN(LI+UIF)
input
   $\mathcal{HC}$  : Horn clauses
vars
   $R$  : resolution tree
   $C$  : conjunctive constraint
   $P$  : proof tree
   $A$  : annotated proof tree
output
   $\Sigma$  : solution
begin
1   $R$  := exhaustively apply RINIT and RSTEP on  $\mathcal{HC}$ 
2   $C$  :=  $\bigwedge$   $leaves(R)$ 
3  if exists  $P$  inferred from  $C$  by PHYP, PCOMB, and PCONG
4    such that  $\models (root(P) \rightarrow 1 \leq 0)$ 
5  then
6     $A$  := exhaustively apply AHYP, ACOMB, and ACONG on  $P$ 
7     $false [ II ]$  :=  $root(A)$ 
8     $\Sigma$  :=  $\{(o, \pi) \mid (o, \pi) \in II \wedge o \notin (leaves(R) \cup \{false\})\}$ 
9    return  $\Sigma$ 
10 else
11 return “no solution exists”
end.

```

Fig. 3. Solving algorithm SOLVEHORN(LI+UIF). Line 7 extracts the partial solution II annotating the root node of A . Line 8 obtains Σ by restricting the domain of II to intermediate nodes of R , i.e., to the nodes that are labeled by queries.

each query q in s with the constraint $\Sigma(q')\sigma$ where q' is in the domain of Σ , queries q' and q have an equal predicate symbol, and $\sigma = match((q', q))$. For example, the above Σ transforms the clause $x \leq f(y) \wedge p(x) \wedge q(y, z) \rightarrow false$ into $x \leq f(y) \wedge (x \leq 0) \wedge (y \leq f(z)) \rightarrow false$.

Let \mathcal{HC}_Σ be the set of query-free clauses obtained by applying Σ . Σ is a *solution* for \mathcal{HC} if each clause c_Σ in \mathcal{HC}_Σ is a valid implication, i.e., $\models c_\Sigma$, and the following condition holds for the uninterpreted function symbols occurring in the range of the solution function. An uninterpreted function symbol f can occur in the solution $\Sigma(q)$ for a query q if f appears in the atoms of a Horn clause from \mathcal{HC} whose head depends on q and in the atoms of a Horn clause from \mathcal{HC} , whose head does not depend on q . For example, given the clauses $\{f(v_1) = 0 \rightarrow p(v_1), f(v_2) = 1 \rightarrow q(v_2), v_3 = v_4 \wedge p(v_3) \wedge q(v_4) \rightarrow false\}$ the function symbol f can appear in the solution of each query. A set of clauses is *satisfiable* if it has a solution.

$$\text{RINIT} \frac{a_1 \wedge \dots \wedge a_m \rightarrow h}{\{(a_1, \dots, a_m, h)\}}$$

$$\text{RSTEP} \frac{\begin{array}{c} R_1 \quad \dots \quad R_n \\ q_1 \wedge \dots \wedge q_n \wedge a_1 \wedge \dots \wedge a_m \rightarrow h \\ R_1 \sigma \cup \dots \cup R_n \sigma \cup \\ \{(q_1, \dots, q_n, a_1, \dots, a_m, h)\} \sigma \end{array}}{\sigma = (\text{match}(q_1, \text{root}(R_1)) + \dots + \text{match}(q_n, \text{root}(R_n)))}$$

Fig. 4. Resolution tree inference rules RINIT and RSTEP.

4 Algorithm

Our goal is an algorithm for computing solutions for recursion-free Horn clauses over linear arithmetic, uninterpreted functions, and queries. This section presents our solving algorithm SOLVEHORN(LI+UIF).

See Figure 3. The algorithm SOLVEHORN(LI+UIF) consists of the following main steps. First, we compute a resolution tree R on the given set of Horn clauses. Next, we take a conjunction C of the leaves of the resolution tree and attempt to find a proof of its unsatisfiability. If no such proof can be found, then we report that there is no solution for the given set of Horn clauses. Otherwise, we proceed with the given proof by annotating its steps. Each intermediate atom derived by proof is annotated by a function that assigns constraints to nodes of the resolution tree. Finally, the annotation of the root of the proof yields a solution for the given set of Horn clauses.

In the rest of this section we provide a detailed presentation of the main steps of SOLVEHORN(LI+UIF).

4.1 Resolution tree

We put together individual Horn clauses from \mathcal{HC} by applying resolution inference. A *resolution tree* keeps the intermediate results of this computation. An edge of a *resolution tree* is a sequence of queries and atoms that is terminated by a query or *false*. Each edge consists of $n > 2$ elements. The first $n - 1$ elements represent the children nodes and the n -th element represents the parent node.

Given the set of Horn clauses \mathcal{HC} , we compute the corresponding resolution tree by applying the inference rules shown in Figure 4. Each rule takes as a premise a set of resolution trees and a Horn clause and infers an extended resolution tree.

The rule RINIT initiates the resolution tree computation by inferring a tree from each clause that does not have any queries in its body. The atoms a_1, \dots, a_m become the children of the node h . The rule RSTEP performs the extension of a set of trees computed so far using a Horn clause. The extension is only possible if the root nodes of the respective trees can be unified with the queries occurring in the body of the clause. This condition is formalized by the side condition requiring the existence of the most general unifier σ . The computed unifier is

$$\begin{array}{c}
\text{PHYP} \frac{}{t \leq 0} t \leq 0 \in \text{atoms}(C) \qquad \text{PCOMB} \frac{t_1 \leq 0 \quad \dots \quad t_n \leq 0}{\lambda_1 t_1 + \dots + \lambda_n t_n \leq 0} \lambda_1, \dots, \lambda_n > 0 \\
\\
\begin{array}{c}
t_1 - s_1 \leq 0 \quad s_1 - t_1 \leq 0 \\
\vdots \qquad \qquad \qquad \vdots \\
t_n - s_n \leq 0 \quad s_n - t_n \leq 0
\end{array} \\
\text{PCONG} \frac{}{f(t_1, \dots, t_n) - f(s_1, \dots, s_n) \leq 0} f(t_1, \dots, t_n), f(s_1, \dots, s_n) \in \text{subterms}(C)
\end{array}$$

Fig. 5. Standard, complete proof rules PHYP, PCOMB, and PCONG for combination of linear rational/real arithmetic and uninterpreted functions. C is the conjunction of leaves of the resolution tree R obtained from the Horn clauses \mathcal{HC} .

applied on the trees and the clause before they are combined into an extended resolution tree.

The resolution tree computation terminates since \mathcal{HC} is recursion-free. Let R be the resulting tree. We consider the set of leaves of the tree, and take their conjunction $C = \bigwedge \text{leaves}(R)$.

For a node o of the resolution tree, we define $\text{insym}(o)$ to be variables and uninterpreted function symbols that occur in atoms in the leaves of the subtree of o , and let $\text{outsym}(o)$ be variables and uninterpreted function symbols that occur in the leaves outside of the subtree of o . Formally, we have

$$\begin{aligned}
\text{insym}(o) &= \bigcup \{ \text{sym}(o') \mid o' \in \text{leaves}(\text{subtree}(o, R)) \} , \\
\text{outsym}(o) &= \bigcup \{ \text{sym}(o') \mid o' \in (\text{leaves}(R) \setminus \text{leaves}(\text{subtree}(o, R))) \} .
\end{aligned}$$

The following proposition allows a transition from the clausal structure to the conjunction of atoms.

Proposition 1. *The set of Horn clauses \mathcal{HC} is satisfiable if and only if the conjunction C is not satisfiable.*

The proof of Proposition 1 follows directly by applying induction over the resolution tree and relying on the definitions of RINIT and RSTEP.

4.2 Proof tree

The algorithm SOLVEHORN(LI+UIF) relies on unsatisfiability proofs. We use a standard set of proof rules for the combination of linear rational/real arithmetic and uninterpreted functions [17]. The implementation of the corresponding proof search procedure is irrelevant for our algorithm, yet we assume that this procedure is complete and use an existing tool for this task, e.g. [4, 7].

See Figure 5 for the proof rules, which we apply to the conjunction of atoms C . The rule PHYP states that atoms appearing in C are provable from C . The rule PCOMB infers that a set of inequalities implies a non-negatively weighted sum thereof. The congruence rule PCONG represents a form of the functionality axiom, which states that equal inputs to a function lead to equal

$$\begin{array}{c}
\text{AHYP} \frac{}{t \leq 0 [\text{MKHYP}(t \leq 0)]} \\
\\
\text{ACOMB} \frac{t_1 \leq 0 [\Pi_1] \quad \dots \quad t_n \leq 0 [\Pi_n]}{\lambda_1 t_1 + \dots + \lambda_n t_n \leq 0 [\text{MKCOMB}(\Pi_1, \dots, \Pi_n, \lambda_1, \dots, \lambda_n)]} \\
\\
\text{ACONG} \frac{
\begin{array}{cc}
t_1 - s_1 \leq 0 [\Pi_1] & s_1 - t_1 \leq 0 [\Pi'_1] \\
\vdots & \vdots \\
t_n - s_n \leq 0 [\Pi_n] & s_n - t_n \leq 0 [\Pi'_n]
\end{array}
}{
\begin{array}{c}
f(t_1, \dots, t_n) - f(s_1, \dots, s_n) \leq 0 [\text{MKCONG}(f(t_1, \dots, t_n), f(s_1, \dots, s_n), \\
\Pi_1, \dots, \Pi_n, \Pi'_1, \dots, \Pi'_n)]
\end{array}
}
\end{array}$$

Fig. 6. Annotation rules. The function MKHYP, MKCOMB, and MKCONG are shown in Figure 7.

results. We are only interested in one inequality part of this axiom. The side condition of PCONG is taken from the interpolating proof rules of [17], and simplifies the proof tree annotation in a way similar to [17].

We assume that there exists a mechanism that uniquely identifies the nodes of the proof tree, even in the presence of nodes that are labeled by equal inequalities, for example by numbering them. For clarity of exposition, we omit any details of such mechanism and assume that the node label carries all necessary information.

If no proof can be found then our algorithm reports that no solution exists. Otherwise, let P be the discovered proof. We assume that P is represented by a tree where nodes are atoms and the children of a node are defined by the rules PHYP, PCOMB, and PCONG. Furthermore, we assume that each edge is labeled by the name of the proof rule that created it.

4.3 Annotated proof tree

We construct a solution for the given Horn clauses through an iterative process, where the intermediate results are called *partial solutions*. Each partial solution is parameterized by a constraint c . A c -partial solution Π for the resolution tree R is a function from nodes of the resolution tree, $nodes(R)$, to constraints that satisfies the following conditions.

$$(\forall o \in leaves(R) : (\models o \rightarrow \Pi(o))) \wedge \quad (\text{PS1})$$

$$(\forall (o^1, \dots, o^m, o) \in R : \models \Pi(o^1) \wedge \dots \wedge \Pi(o^m) \rightarrow \Pi(o)) \wedge \quad (\text{PS2})$$

$$(\models \Pi(false) \rightarrow c) \wedge \quad (\text{PS3})$$

$$(\forall o \in nodes(R) : sym(\Pi(o)) \subseteq (insym(o) \cap outsym(o)) \cup sym(c)) \quad (\text{PS4})$$

Our annotation uses constraints of the following form, called *solution constraints*.

$$\text{solution constraints } \ni \pi ::= t \leq 0 \mid \hat{c} \wedge (\hat{c} \rightarrow \pi)$$

To simplify the presentation, we represent a solution constraint

$$C_1 \wedge (D_1 \rightarrow (\dots C_r \wedge (D_r \rightarrow p \leq 0)))$$

as a pair consisting of a corresponding sequence and a term $\langle ((C_1, D_1), \dots, (C_r, D_r)), p \rangle$. A solution constraint $p \leq 0$ is represented by $\langle [], p \rangle$.

Given the proof tree P , we annotate its nodes with partial solutions using the rules shown in Figure 6 and auxiliary functions shown in Figure 7. The rule AHYP annotates each leaf of the proof tree with the result of applying the function MKHYP. The annotation is enclosed by a pair of square brackets. The rule ACOMB shows how to annotate a parent node when provided with an annotation of its children in case when the parent was obtained by a non-negatively weighted sum. The parent annotation is computed by MKCOMB. Similarly, the rule ACONG annotates parent nodes obtained by the congruence rule.

For each node of R at line 6, ACONG has four cases that deal with the difficulty of solving Horn clauses over uninterpreted functions, i.e., a sub term may contain variables that are not allowed to appear in the partial solutions. The proof of theorem 3 explains how these cases avoid such variables in the partial solutions.

We annotate P and obtain an annotated proof tree A . Our algorithm SOLVE-HORN(LI+UIF) uses the annotation of the root of A to derive a solution to the Horn clauses \mathcal{HC} .

5 Correctness and complexity

This section presents the correctness and complexity properties of our algorithm. The corresponding proofs are in Appendix C of extended version of this paper [10].

The correctness of our algorithm follows from Proposition 1 and Theorems 1–3 below. First, we establish that a $(1 \leq 0)$ -partial solution, which satisfies Equations (PS1)–(PS4), defines a solution for the given Horn clauses.

Theorem 1. *$(1 \leq 0)$ -partial solution defines a solution of the Horn clauses.*

Now, we show that the annotations computed by the rules in Figure 6 satisfy the partial solution conditions in Equations (PS1)–(PS4). This step relies on the following inductive invariant.

Definition 1 (*$t \leq 0$ -annotation invariant*). *Π is $t \leq 0$ -annotation invariant for the resolution tree R if there exists $r \geq 0$ such that for each $o \in \text{nodes}(R)$ the following conditions hold.*

- $\Pi(o)$ is a solution constraint such that

$$\Pi(o) = \langle ((C_1, D_1), \dots, (C_r, D_r)), p \rangle. \quad (\text{AI-1})$$

```

function MKHYP
input
   $t \leq 0$  : inequality term/node in  $R$ 
begin
1 for each  $o \in \text{nodes}(R)$  do
2   if  $t \leq 0 \in \text{leaves}(\text{subtree}(o, R))$  then
3      $\Pi(o) := \langle \square, t \rangle$ 
4   else
5      $\Pi(o) := \langle \square, 0 \rangle$ 
6   return  $\Pi$ 
7 end

function MKCOMB
input
   $\Pi_1, \dots, \Pi_n$  : partial solutions
   $\lambda_1, \dots, \lambda_n$  : constants
begin
1 for each  $o \in \text{nodes}(R)$  do
2   for each  $i \in 1..n$  do
3      $\langle L_i, t_i \rangle := \Pi_i(o)$ 
4      $L := L_1 \bullet \dots \bullet L_n$ 
5      $t := \lambda_1 t_1 + \dots + \lambda_n t_n$ 
6      $\Pi(o) := \langle L, t \rangle$ 
7   return  $\Pi$ 
8 end

function MKCONG
input
   $f(t_1, \dots, t_n), f(s_1, \dots, s_n)$  : terms
   $\Pi_1, \dots, \Pi_n, \Pi'_1, \dots, \Pi'_n$  : partial solutions
begin
1 for each  $o \in \text{nodes}(R)$  do
2   for each  $i \in 1..n$  do
3      $\langle L_i, p_i \rangle := \Pi_i(o)$ 
4      $\langle L'_i, p'_i \rangle := \Pi'_i(o)$ 
5      $(C, D, p) :=$ 
6       match  $\text{sym}(f(t_1, \dots, t_n)) \subseteq \text{outsym}(o),$ 
7          $\text{sym}(f(s_1, \dots, s_n)) \subseteq \text{outsym}(o)$  with
8       |  $\text{true}, \text{true} \rightarrow (\bigwedge_{i=1}^n (p_i \leq 0 \wedge p'_i \leq 0), \text{true}, 0)$ 
9       |  $\text{true}, \text{false} \rightarrow (\bigwedge_{i=1}^n p_i + p'_i \leq 0, \bigwedge_{i=1}^n -p_i - p'_i \leq 0,$ 
10          $f(s_1 + p_1, \dots, s_n + p_n) - f(s_1, \dots, s_n))$ 
11       |  $\text{false}, \text{true} \rightarrow (\bigwedge_{i=1}^n p_i + p'_i \leq 0, \bigwedge_{i=1}^n -p_i - p'_i \leq 0,$ 
12          $f(t_1, \dots, t_n) - f(t_1 + p'_1, \dots, t_n + p'_n))$ 
13       |  $\text{false}, \text{false} \rightarrow (\text{true}, \bigwedge_{i=1}^n (t_i - s_i - p_i \leq 0 \wedge s_i - t_i - p'_i \leq 0),$ 
14          $f(t_1, \dots, t_n) - f(s_1, \dots, s_n))$ 
15      $\Pi(o) := \langle L_1 \bullet \dots \bullet L_n \bullet L'_1 \bullet \dots \bullet L'_n \bullet (C, D), p \rangle$ 
16 return  $\Pi$ 
17 end

```

Fig. 7. Computation of partial solutions to annotate nodes of the proof tree, as shown in Figure 6. We use \bullet to denote concatenation of sequences.

– If $o \in \text{leaves}(R)$ then

$$\left(\forall i \in 1..r : \models o \wedge \bigwedge_{k=1}^{i-1} D_k \rightarrow C_i \right) \wedge \quad (\text{AI-2a})$$

$$\left(\models o \wedge \bigwedge_{k=1}^r D_k \rightarrow p \leq 0 \right). \quad (\text{AI-2b})$$

– If $(o^1, \dots, o^m, o) \in R$ and $\forall j \in 1..m : \Pi(o^j) = \langle (C_1^j, D_1^j), \dots, (C_r^j, D_r^j) \rangle, p^j$ then

$$\left(\forall i \in 1..r : \models \left(\bigwedge_{k=1}^i \bigwedge_{l=1}^m C_k^l \right) \wedge \bigwedge_{k=1}^{i-1} D_k \rightarrow C_i \right) \wedge \quad (\text{AI-3a})$$

$$\left(\begin{array}{l} \forall i \in 1..r \\ \forall j \in 1..m \end{array} : \models \left(\bigwedge_{l \in 1..m \setminus \{j\}} C_i^l \right) \wedge \left(\bigwedge_{k=1}^{i-1} \bigwedge_{l=1}^m C_k^l \right) \wedge \bigwedge_{k=1}^i D_k \rightarrow D_i^j \right) \wedge \quad (\text{AI-3b})$$

$$\left(\models \left(\bigwedge_{k=1}^r \bigwedge_{l=1}^m C_k^l \right) \wedge \bigwedge_{k=1}^r D_k \rightarrow p - p^1 - \dots - p^m \leq 0 \right). \quad (\text{AI-3c})$$

– If $o = \text{false}$ then

$$p = t \wedge \forall i \in 1..r : D_i = C_i = \text{true}. \quad (\text{AI-4})$$

– Conditions on symbol appearance:

$$\text{sym}(\{C_1, \dots, C_r, D_1, \dots, D_r, p \leq 0\}) \subseteq \text{insym}(o) \wedge \quad (\text{AI-5})$$

$$\text{sym}(\{C_1, \dots, C_r, D_1, \dots, D_r, t - p \leq 0\}) \subseteq \text{outsym}(o). \quad (\text{AI-6})$$

The above definition act as an intermediate step. In theorem 2, we show that a $t \leq 0$ -annotation invariant satisfies all the conditions for being a $t \leq 0$ -partial solution.

Theorem 2. *Each $t \leq 0$ -annotation invariant is a $t \leq 0$ -partial solution.*

Now, we show that the presented algorithm computes the partial solutions that satisfies the invariant.

Theorem 3. *The annotation rules in Figure 6 compute annotation invariants.*

Theorem 4 (Complexity). *The application of the annotation rules from Figure 6 takes time proportional to the product of the size of the proof tree and the size of the resolution tree. The size of the resolution tree is linear in the size of the corresponding set of recursion-free, tree-like Horn clauses.*

Note that we present the complexity of our algorithm in terms of the size of the proof tree. Since the size of a resolution tree can also be exponential in the size of the set of Horn clauses, the size of a proof tree can be exponential.

6 Conclusion

We presented an algorithm for computing solutions for recursion-free Horn clauses over the combination of linear rational/real arithmetic, uninterpreted functions, and queries.

Connection to interpolation The interpolation algorithm presented in [17] is a special case for the algorithm presented in this paper. As illustrated in the introduction, an interpolation problem can be reduced to solving a set of recursion-free Horn clauses. The set of Horn clauses resulting from an interpolation problem has only one unknown query. Therefore, the corresponding resolution tree obtained from the set of Horn clauses contains only one internal node. The partial solution of this internal node in the $(1 \leq 0)$ -partial solution will be the interpolant. In this special case, we only need to track partial solutions of the internal node in the annotated proof tree. We can transform our algorithm for this case such that nodes of the proof tree are annotated with a formula corresponding to the partial solution of this internal node. The resulting algorithm will be the algorithm presented in [17].

Our algorithm can be directly applied to support abstraction and refinement tasks for the verification of programs with procedures, threads and higher order functions.

7 Acknowledgment

Ashutosh Gupta was supported in part by the DFG Graduiertenkolleg 1480 (PUMA), FWF NFN Grant No S11407-N23 (RiSE), and the ERC Advanced Grant QUAREM.

References

1. D. Beyer, A. Cimatti, A. Griggio, M. E. Keremoglu, and R. Sebastiani. Software model checking via large-block encoding. In *FMCAD*, 2009.
2. D. Beyer, D. Zufferey, and R. Majumdar. CSIsat: Interpolation for LA+EUF. In *CAV*, 2008.
3. A. Brillout, D. Kroening, P. Rümmer, and T. Wahl. An interpolating sequent calculus for quantifier-free Presburger arithmetic. In *Proceedings of IJCAR*, LNCS, pages 384–399. Springer, 2010.
4. R. Bruttomesso, A. Cimatti, A. Franzén, A. Griggio, and R. Sebastiani. The MathSAT 4SMT solver. In *CAV*, 2008.
5. A. Cimatti, A. Griggio, and R. Sebastiani. Interpolant generation for UTVPI. In *CADE*, 2009.
6. A. Cimatti, A. Griggio, and R. Sebastiani. Efficient generation of Craig interpolants in satisfiability modulo theories. *ACM Trans. Comput. Logic*, 12, November 2010.
7. L. M. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, 2008.
8. A. Goel, S. Krstic, and C. Tinelli. Ground interpolation for combined theories. In *CADE*, 2009.

9. A. Gupta, C. Popeea, and A. Rybalchenko. Predicate abstraction and refinement for verifying multi-threaded programs. In *POPL*, 2011.
10. A. Gupta, C. Popeea, and A. Rybalchenko. Solving recursion-free Horn clauses over LI+UIF. available from <http://pub.ist.ac.at/~agupta/papers/HornLIUIF.pdf>, 2011.
11. M. Heizmann, J. Hoenicke, and A. Podelski. Nested interpolants. In *POPL*, 2010.
12. T. A. Henzinger, R. Jhala, R. Majumdar, and K. L. McMillan. Abstractions from proofs. In *POPL*, 2004.
13. H. Jain, E. M. Clarke, and O. Grumberg. Efficient Craig interpolation for linear Diophantine (dis)equations and linear modular equations. *Formal Methods in System Design*, pages 6–39, 2009.
14. R. Jhala and R. Majumdar. Counterexample refinement for functional programs. available from <http://www.cs.ucla.edu/~rupak/Papers/CEGARFunctional.ps>, 2009.
15. R. Jhala and K. L. McMillan. A practical and complete approach to predicate refinement. In *TACAS*, 2006.
16. D. Kroening, J. Leroux, and P. Rümmer. Interpolating quantifier-free Presburger arithmetic. In *Proceedings of LPAR*, LNCS, pages 489–503. Springer, 2010.
17. K. L. McMillan. An interpolating theorem prover. *Theor. Comput. Sci.*, 345(1):101–121, 2005.
18. K. L. McMillan. Lazy abstraction with interpolants. In *CAV*, pages 123–136, 2006.
19. A. Rybalchenko and V. Sofronie-Stokkermans. Constraint solving for interpolation. In *VMCAI*, 2007.
20. T. Terauchi. Dependent types from counterexamples. In *POPL*, 2010.
21. H. Unno and N. Kobayashi. Dependent type inference with interpolants. In *PPDP*, 2009.
22. G. Yorsh and M. Musuvathi. A combination method for generating interpolants. In *CADE*, pages 353–368, 2005.