

Solving Existentially Quantified Horn Clauses

Tewodros A. Beyene¹, Corneliu Popeea¹, and Andrey Rybalchenko^{1,2}

¹ Technische Universität München

² Microsoft Research Cambridge

Abstract. Temporal verification of universal (i.e., valid for all computation paths) properties of various kinds of programs, e.g., procedural, multi-threaded, or functional, can be reduced to finding solutions for equations in form of universally quantified Horn clauses extended with well-foundedness conditions. Dealing with existential properties (e.g., whether there exists a particular computation path), however, requires solving forall-exists quantified Horn clauses, where the conclusion part of some clauses contains existentially quantified variables. For example, a deductive approach to CTL verification reduces to solving such clauses. In this paper we present a method for solving forall-exists quantified Horn clauses extended with well-foundedness conditions. Our method is based on a counterexample-guided abstraction refinement scheme to discover witnesses for existentially quantified variables. We also present an application of our solving method to automation of CTL verification of software, as well as its experimental evaluation.

1 Introduction

Temporal verification of universal, i.e., valid for all computation paths, properties of various kinds of programs is a success story. Various techniques, e.g., abstract domains [13], predicate abstraction [18,22], or interpolation [26], provide a basis for efficient tools for the verification of such properties, e.g., Astree [5], Blast [22], CPAchecker [3], SatAbs [9], Slam [2], Terminator [12], or UFO [1]. To a large extent, the success of checkers of universal properties is determined by tremendous advances in the state-of-the-art in decision procedures for (universal) validity checking, i.e., advent of tools like MathSAT [6] or Z3 [15].

In contrast, advances in dealing with existential properties of programs, e.g., proving whether there exists a particular computation path, are still not on par with the maturity of verifiers for universal properties. Nevertheless, important first steps were made in proving existence of infinite program computations, see e.g. [16, 20, 29], even in proving existential (as well as universal) CTL properties [11]. Moreover, bounded model checking tools like CBMC [8] or Klee [7] can be very effective in proving existential reachability properties. All these initial achievements inspire further, much needed research on the topic.

In this paper, we present a method that can serve as a further building block for the verification of temporal existential (and universal) properties of programs. Our method solves forall-exists quantified Horn clauses extended with

well-foundedness conditions. (The conclusion part of such clauses may contain existentially quantified variables.) The main motivation for the development of our method stems from an observation that verification conditions for existential temporal properties, e.g., generated by a deductive proof system for CTL [25], can be expressed by clauses in such form.

Our method, called E-HSF, applies a counterexample-guided refinement scheme to discover witnesses for existentially quantified variables. The refinement loop collects a global constraint that declaratively determines which witnesses can be chosen. The chosen witnesses are used to replace existential quantification, and then the resulting universally quantified clauses are passed to a solver for such clauses. At this step, we can benefit from emergent tools in the area of solving Horn clauses over decidable theories, e.g., HSF [19], μZ [23], or Duality [27]. Such a solver either finds a solution, i.e., a model for uninterpreted relations constrained by the clauses, or returns a counterexample, which is a resolution tree (or DAG) representing a contradiction. E-HSF turns the counterexample into an additional constraint on the set of witness candidates, and continues with the next iteration of the refinement loop. Notably, our refinement loop conjoins constraints that are obtained for all discovered counterexamples. This way E-HSF guarantees that previously handled counterexamples are not rediscovered and that a wrong choice of witnesses can be mended.

We applied our implementation of E-HSF to forall-exists quantified Horn clauses with well-foundedness conditions that we obtained by from a deductive proof system for CTL [25]. The experimental evaluation on benchmarks from [11] demonstrates the feasibility of our method.

2 Preliminaries

In this section we introduce preliminary definitions.

Constraints Let \mathcal{T} be a first-order theory in a given signature and $\models_{\mathcal{T}}$ be the entailment relation for \mathcal{T} . We write v, v_0, v_1, \dots and w to denote non-empty tuples of variables. We refer to a formula $c(v)$ over variables v from \mathcal{T} as a constraint. Let *false* and *true* be an unsatisfiable and a valid constraint, respectively.

For example, let x, y , and z be variables. Then, $v = (x, y)$ and $w = (y, z)$ are tuples of variables. $x \leq 2$, $y \leq 1 \wedge x - y \leq 0$, and $f(x) + g(x, y) \leq 3 \vee z \leq 0$ are example constraints in the theory \mathcal{T} of linear inequalities and uninterpreted functions, where f and g are uninterpreted function symbols. $y \leq 1 \wedge x - y \leq 0 \models_{\mathcal{T}} x \leq 2$ is an example of a valid entailment.

A binary relation is well-founded if it does not admit any infinite chains. A relation $\varphi(v, v')$ is disjunctively well-founded if it is included in a finite union of well-founded relations [31], i.e., if there exist well-founded $\varphi_1(v, v'), \dots, \varphi_n(v, v')$ such that $\varphi(v, v') \models_{\mathcal{T}} \varphi_1(v, v') \vee \dots \vee \varphi_n(v, v')$. For example, the relation $x \geq 0 \wedge x' \leq x - 1$ is well-founded, while the relation $(x \geq 0 \wedge x' \leq x - 1) \vee (y \leq 0 \wedge y' \geq y + 1)$ is disjunctively well-founded.

Queries and dwf-predicates We assume a set of uninterpreted predicate symbols \mathcal{Q} that we refer to as query symbols. The arity of a query symbol is encoded in its name. We write q to denote a query symbol. Given q of a non-zero arity n and a tuple of variables v of length n , we define $q(v)$ to be a query. Furthermore, we introduce an interpreted predicate symbol dwf of arity one (dwf stands for disjunctive well-foundedness). Given a query $q(v, v')$ over tuples of variables with equal length, we refer to $dwf(q)$ as a dwf -predicate. For example, let $\mathcal{Q} = \{r, s\}$ be query symbols of arity one and two, respectively. Then, $r(x)$ and $s(x, y)$ are queries, and $dwf(s)$ is a dwf -predicate.

Forall-exists Horn-like clauses Let $h(v)$ range over queries over v , constraints over v , and existentially quantified conjunctions of queries and constraints with free variables in v . We define a forall-exists Horn-like clause to be either an implication $c(v_0) \wedge q_1(v_1) \wedge \dots \wedge q_n(v_n) \rightarrow h(v)$ or a unit clause $dwf(q)$. The left-hand side of the implication is called the body, written as $body(v)$, and the right-hand side is called the head.

We give as example a set of forall-exists Horn-like clauses below:

$$\begin{aligned} x \geq 0 \rightarrow \exists y : x \geq y \wedge rank(x, y), & \quad rank(x, y) \rightarrow ti(x, y), \\ ti(x, y) \wedge rank(y, z) \rightarrow ti(x, z), & \quad dwf(ti). \end{aligned}$$

These clauses represent an assertion over the interpretation of predicate symbols $rank$ and ti .

Semantics of forall-exists Horn-like clauses A set of clauses can be seen as an assertion over the queries that occur in the clauses.

We consider a function $ClauseSol$ that maps each query $q(v)$ occurring in a given set of clauses into a constraint over v . Such a function is called a solution if the following two conditions hold. First, for each clause of the form $body(v) \rightarrow h(v)$ from the given set we require that replacing each query by the corresponding constraint assigned by $ClauseSol$ results in a valid entailment. That is, we require $body(v) ClauseSol \models_{\mathcal{T}} h(v) ClauseSol$, where the juxtaposition represents application of substitution. Second, for each clause of the form $dwf(q)$ we require that the constraint assigned by $ClauseSol$ to q represents a disjunctively well-founded relation. Let $\models_{\mathcal{Q}}$ be the corresponding satisfaction relation, i.e., $ClauseSol \models_{\mathcal{Q}} Clauses$ if $ClauseSol$ is a solution for the given set of clauses.

For example, the previously presented set of clauses, say $Clauses$, has a solution $ClauseSol$ such that $ClauseSol(rank(x, y)) = ClauseSol(ti(x, y)) = (x \geq 0 \wedge y \geq x - 1)$. To check $ClauseSol \models_{\mathcal{Q}} Clauses$ we consider the validity of the following implications:

$$\begin{aligned} x \geq 0 \rightarrow \exists y : x \geq y \wedge x \geq 0 \wedge y \leq x - 1, \\ x \geq 0 \wedge y \leq x - 1 \rightarrow x \geq 0 \wedge y \leq x - 1, \\ x \geq 0 \wedge y \leq x - 1 \wedge y \geq 0 \wedge z \leq y - 1 \rightarrow x \geq 0 \wedge z \leq x - 1. \end{aligned}$$

and the fact that $ClauseSol(ti(x, y)) = (x \geq 0 \wedge y \leq x - 1)$ is a (disjunctively) well-founded relation.

```

function SKOLEMIZE(Clauses)
1  Skolemized := Parent := Rels := Grds :=  $\emptyset$ 
2  for each clause  $\in$  Clauses do
3    match clause with
4    | body(v)  $\rightarrow \exists w : \bigwedge_{i=1}^n \text{conj}_i(v, w) \rightarrow$ 
5      rel, grd := fresh predicate symbols of arity  $|v| + |w|$  and  $|v|$ , resp.
6      Parent :=  $\{(grd, clause), (rel, clause)\} \cup Parent$ 
7      Rels :=  $\{rel\} \cup Rels$ 
8      Grds :=  $\{grd\} \cup Grds$ 
9      Skolemized :=  $\{body(v) \wedge rel(v, w) \rightarrow conj_i(v, w) \mid i \in 1..n\} \cup$ 
10          $\{body(v) \rightarrow grd(v)\} \cup Skolemized$ 
11    |  $\_ \rightarrow Skolemized$  :=  $\{clause\} \cup Skolemized$ 
12  done
13  return (Skolemized, Parent, Rels, Grds)

```

Fig. 1. Function SKOLEMIZE replaces existential quantification by application of Skolem relations. In line 4, each $\text{conj}_i(v, w)$ is either a query or a constraint.

Solving Horn-like clauses without existential quantification We assume an algorithm HSF for solving Horn-like clauses whose heads do not contain any existential quantification. This algorithm computes a solution *ClauseSol* when it exists. There already exist such algorithms as well as their efficient implementations that are based on predicate abstraction and interpolation [19], as well as interpolation based approximation [27].

3 Solving algorithm E-HSF

In this section we present our algorithm E-HSF for solving constraints in form of Horn clauses that contain existential quantification and well-foundedness conditions.

Our solving method proceeds in two steps. First, we rely on Skolemization to re-formulate the problem of dealing with existential quantification as a problem of finding witnesses for the existentially quantified variables. Such witnesses are represented by Skolem relations (which is a slight generalisation of Skolem functions that is convenient in our setting). For an existentially quantified clause $body(v) \rightarrow \exists w : head(v, w)$, the corresponding Skolem relation $rel(v, w)$ determines which value w satisfies $head(v, w)$ for a given v . Since for each v such that $body(v)$ holds we need a value w , we require that such v is in the domain of the Skolem relation. We represent the domain of Skolem relation $rel(v, w)$ as the guard $grd(v)$, and will use it later to implement the above requirement.

A function SKOLEMIZE shown in Figure 1 implements the Skolemization step. It outputs a set of clauses without existential quantification, yet containing Skolem relations and guards. Furthermore, SKOLEMIZE keeps track of which Skolem relations and guards belong to which clauses.

The second step takes as input a set of Skolemized clauses produced by SKOLEMIZE and either finds a solution, returns that no solution can be found, or diverges. At this step we rely on a set of templates that determine the search space for Skolem relations, their guards, as well as termination arguments used for dealing with well-foundedness. In order to ensure that the guard of a Skolem relation entails its domain, we assume that the guard template implies the projection of the Skolem relation template. Formally, we require that the template functions GRDT and RELT providing guard and Skolem relation templates for the output of SKOLEMIZE satisfy the following condition: for each $grd \in Grds$ and $rel \in Rels$ such that $Parent(grd) = Parent(rel)$ the implication

$$GRDT(grd)(v) \rightarrow \exists w : RELT(rel)(v, w) \quad (1)$$

is valid (for arbitrary values of template parameters). We establish Equation 1 by choosing templates accordingly.

See Figure 2. The solving process iteratively determines appropriate candidates for Skolem relations and their guards by using a counterexample driven approach. Each counterexample induces constraints on template parameters and thus rules out failed attempts. Given candidates for Skolem relations and their guards, we record these candidates by introducing appropriate Horn clauses called *Defs*. Then, we apply a solver for (ordinary) Horn clauses, which we call HSF, on the set of Skolemized clauses that is extended with *Defs*. If HSF finds a solution, then we report it as a solution for the original set of clauses. Otherwise, we inspect a counterexample given by HSF. Such a counterexample is presented by a set of recursion-free Horn clauses which uses a form of Static Single Assignment (SSA) to represent an unfolding of $Skolemized \cup Defs$ that cannot be satisfied.

If the counterexample does not involve any Skolem relations or their guards, then we report that *Skolemized* cannot be satisfied. Otherwise, the unfolding is not satisfiable either because there are no Skolem relations together with guards that make *Skolemized* satisfiable, or because the currently chosen candidates are not correct. To find out, we replace the candidates by templates and create a constraint over template parameters ensuring that the counterexample is eliminated. This constraint is determined by ENCODEVALIDITY and results in a formula whose free variables are template parameters. We consider a conjunction of such constraints, which is stored as *Constraint*, across all iteration of the solving process, thus ensuring that previously analysed and eliminated counterexamples will not re-appear. A solution of *Constraint* determines new candidates, which we formally record using the set of clauses *Defs*. Now our iteration is ready to go in the next round.

Correctness The algorithm E-HSF relies on the following propositions. First, the Skolemization step preserves equi-satisfiability under an assumption that each guard needs to be a subset of the corresponding Skolem relation.

Lemma 1 (Skolemization preserves satisfiability). The set of clauses *Clauses* is equi-satisfiable with the set of clauses computed by SKOLEMIZE when

```

algorithm E-HSF(Clauses)
1  Skolemized, Parent, Rels, Grds := SKOLEMIZE(Clauses)
2  Constraint := true
3  Defs := {true → rel(v, w) | rel ∈ Rels} ∪ {grd(v) → true | grd ∈ Grds}
4  match HSF(Skolemized ∪ Defs) with
5  | solution ClauseSol → return “solution ClauseSol”
6  | error derivation Cex and symbol map SYM →
7    CexDefs := {(body → q(...)) ∈ Cex | SYM(q) ∈ Rels ∪ Grds}
8    if CexDefs = ∅ then return “error derivation Cex and symbol map SYM”
9    else
10     (body ∧ ∧i=1n qi(vi, wi) → head) := RESOLVE(Cex \ CexDefs)
11     body := body ∧ ∧i=1n RELT(SYM(qi))(vi, wi)
12     match head with
13     | q(v, w) when dwf(SYM(q)) ∈ Clauses →
14       head := BOUND(SYM(q))(v) ∧ DECREASET(SYM(q))(v, w)
15     | q(v) when SYM(q) ∈ Grds →
16       head := GRDT(SYM(q))(v)
17     | _ → skip
18     Constraint := ENCODEVALIDITY(body → head) ∧ Constraint
19     match SMTSOLVE(Constraint) with
20     | solution CexSol →
21       Defs := {RELT(rel)(v, w) CexSol → rel(v, w) | rel ∈ Rels} ∪
22         {grd(v) → GRDT(grd)(v) CexSol | grd ∈ Grds}
23     goto line 4
24     | _ → return “error derivation Cex and symbol map SYM”

```

Fig. 2. Solving algorithm E-HSF for Horn clauses with existential quantification and disjunctive well-foundedness predicates. RESOLVE applies resolution. ENCODEVALIDITY encodes the Farkas’ lemma from linear programming [32], while SMTSOLVE returns an assignment of constants to template parameters provided its argument is satisfiable.

domains of Skolem relations contain corresponding guards. Formally, *Clauses* is equi-satisfiable with the set

$$\{grd(v) \rightarrow \exists w : rel(v, w) \mid grd \in Grds \wedge rel \in Rels \wedge Parent(grd) = Parent(rel)\} \cup Skolemized .$$

Proof. (Sketch) Let *clause* = (*body*(*v*) → ∃*w* : *q*(*v*, *w*)) and *Parent*(*rel*) = *Parent*(*grd*) = *clause*. We keep pairs (*v*^{*}, *w*^{*}) such that *body*(*v*^{*}) and *q*(*v*^{*}, *w*^{*}) hold in a relation *rel* while storing *v*^{*} in *grd*. Then the statement of the lemma follows immediately. □

As previously mentioned, the above relation between Skolem relations and their guards is established by the appropriate choice of RELT and GRDT functions, see Equation 1. Then E-HSF inherits its soundness from HSF.

Theorem 1 (Soundness). If HSF is sound, i.e., it returns solutions for given sets of clauses, and if Equation 1 holds for each $grd \in Grds$ and $rel \in Rels$ such that $Parent(grd) = Parent(rel)$, then, upon termination, E-HSF returns a solution for *Clauses*.

Proof. Let *ClauseSol* be a result of applying HSF in line 5 of Figure 2. The first assumption of the theorem statement guarantees that *ClauseSol* satisfies *Skolemized*. The first assumption ensures that Lemma 1 is applicable, hence, *ClauseSol* satisfies *Clauses*. \square

Our method is based on a counterexample guided scheme for discovery of Skolem relations and guards. While this scheme has successful applications in practice, it does not guarantee termination of the refinement process when the set of candidates for Skolem relations and guards is unbounded. Our method necessarily inherits this limitation.

Despite this undecidability imposed limitation, our method strives at achieving termination of refinement process in practice. An important ingredient is provided by the fact that *Constraint* keeps track of the conjunction of constraints used to discover candidates Skolem relations and guards across all iterations.

Theorem 2 (Progress of refinement). E-HSF does not consider any error derivation/counterexample more than once.

Proof. (Sketch) The progress of refinement property follows directly from the observation that every solution for *Constraint* yields Skolem relations and guards that satisfy each previously discovered error derivation. \square

4 Example of applying E-HSF

We consider the following set *Clauses* that encodes a check whether a program with the variables $v = (x, y)$, an initial condition $init(v) = (y \geq 1)$ and a transition relation $next(v, v') = (x' = x + y)$ satisfies a CTL property $EF\ dst(v)$, where $dst(v) = (x \geq 0)$.

$$\begin{aligned} init(v) &\rightarrow inv(v), & inv(v) \wedge \neg dst(v) &\rightarrow \exists v' : next(v, v') \wedge inv(v') \wedge rank(v, v'), \\ rank(v, v') &\rightarrow ti(v, v'), & ti(v, v') \wedge rank(v', v'') &\rightarrow ti(v, v''), & dwf(ti). \end{aligned}$$

Here, $inv(v)$, $rank(v, v')$, and $ti(v, v')$ are unknown predicates that we need to solve for. The predicate $inv(v)$ corresponds to states reachable during program execution, while the second row of clauses ensures that $rank(v, v')$ is a well-founded relation [31].

We start the execution of E-HSF from Figure 2 by applying SKOLEMIZE to eliminate the existential quantification. As a result, the clause that contains existential quantification is replaced by the following four clauses that contain an application of a Skolem relation $rel(v, v')$ introduced by SKOLEMIZE as well as an introduction of a lower bound on the guard $grd(v)$ of the Skolem relation:

$$\begin{aligned} inv(v) \wedge \neg dst(v) \wedge rel(v, v') &\rightarrow next(v, v'), \\ inv(v) \wedge \neg dst(v) \wedge rel(v, v') &\rightarrow inv(v'), \\ inv(v) \wedge \neg dst(v) \wedge rel(v, v') &\rightarrow rank(v, v'), \\ inv(v) \wedge \neg dst(v) &\rightarrow grd(v). \end{aligned}$$

Furthermore, this introduction is recorded as $Rels = \{rel\}$ and $Grds = \{grd\}$. Note that we replaced a conjunction in the head of a clause by a conjunction of corresponding clauses.

First candidate for Skolem relation Next, we proceed with the execution of E-HSF. We initialise *Constraint* with the assertion *true*. Then, we generate a set of Horn clauses *Defs* that provides initial candidates for the Skolem relation and its guard as follows: $Defs = \{true \rightarrow rel(v, v'), grd(v) \rightarrow true\}$. Now, we apply the solving algorithm HSF for quantifier free Horn clauses on the set of clauses that contains the result of Skolemization and the initial candidates in *Defs*, i.e., we give to HSF the following clauses:

$$\begin{aligned} init(v) &\rightarrow inv(v), & rank(v, v') &\rightarrow ti(v, v'), \\ inv(v) \wedge \neg dst(v) \wedge rel(v, v') &\rightarrow next(v, v'), & ti(v, v') \wedge rank(v', v'') &\rightarrow ti(v, v''), \\ inv(v) \wedge \neg dst(v) \wedge rel(v, v') &\rightarrow inv(v'), & dwf(ti), \\ inv(v) \wedge \neg dst(v) \wedge rel(v, v') &\rightarrow rank(v, v'), & true &\rightarrow rel(v, v'), \\ inv(v) \wedge \neg dst(v) &\rightarrow grd(v), & grd(v) &\rightarrow true. \end{aligned}$$

HSF returns an error derivation that witnesses a violation of the given set of clauses. This derivation represents an unfolding of clauses in $Skolemized \cup Defs$ that yields a relation for $ti(v, v')$ that is not disjunctively well-founded. To represent the unfolding, HSF uses a form of static single assignment (SSA) that is applied to predicate symbols, where each unfolding step introduces a fresh predicate symbol that is recorded by the function SYM. We obtain the clauses *Cex* consisting of

$$init(v) \rightarrow q_1(v), \quad q_1(v) \wedge \neg dst(v) \wedge q_2(v, v') \rightarrow next(v, v'), \quad true \rightarrow q_2(v, v')$$

together with the following bookkeeping of the SSA renaming: $SYM(q_1) = inv$ and $SYM(q_2) = rel$. From *Cex* we extract the clause *CexDefs* that provides the candidate for the Skolem relation. We obtain $CexDefs = \{true \rightarrow q_2(v, v')\}$, since $SYM(q_2) = rel$ and hence $SYM(q_2) \in Rels$.

We analyse the counterexample clauses by applying resolution on $Cex \setminus CexDefs$. The corresponding resolution tree is shown below (literals

selected for resolution are boxed):

$$\frac{\text{init}(v) \rightarrow \boxed{q_1(v)} \quad \boxed{q_1(v)} \wedge \neg \text{dst}(v) \wedge q_2(v, v') \rightarrow \text{next}(v, v')}{\text{init}(v) \wedge \neg \text{dst}(v) \wedge q_2(v, v') \rightarrow \text{next}(v, v')}$$

Note that $q_2(v, v')$ was not eliminated, since the clause $\text{true} \rightarrow q_2(v, v')$ was not given to RESOLVE as input. The result of applying RESOLVE is the clause $\text{init}(v) \wedge \neg \text{dst}(v) \wedge q_2(v, v') \rightarrow \text{next}(v, v')$. We assign the conjunction $\text{init}(v) \wedge \neg \text{dst}(v)$ to *body* and $\text{next}(v, v')$ to *head*, respectively.

Now we iterate i through the singleton set $\{1\}$, which is determined by the fact that the above clause contains only one unknown predicate on the left-hand side. We apply RELT on $\text{SYM}(q_2)$ and set the free variables in the result to (v, v') . This yields a template $v' = Tv + t$ for the Skolem relation $\text{rel}(v, v')$. Here, T is a matrix of unknown coefficients $\begin{pmatrix} t_{xx} & t_{xy} \\ t_{yx} & t_{yy} \end{pmatrix}$, and t is a vector of unknown free coefficient (t_x, t_y) . In other words, our template represents a conjunction of two equality predicates $x' = t_{xx}x + t_{xy}y + t_x$ and $y' = t_{yx}x + t_{yy}y + t_y$. We conjoin this template with *body* and obtain $\text{body} = (v' = Tv + t \wedge \text{init}(v) \wedge \neg \text{dst}(v))$. Since *head* is not required to be disjunctively well-founded, E-HSF proceeds with the generation of constraints over template parameters.

We apply ENCODEVALIDITY on the following implication:

$$x' = t_{xx}x + t_{xy}y + t_x \wedge y' = t_{yx}x + t_{yy}y + t_y \wedge y \geq 1 \wedge \neg x \geq 0 \rightarrow x' = x + y .$$

This implication is valid if the following constraint returned by ENCODEVALIDITY is satisfiable.

$$\exists \overbrace{\lambda_1, \lambda_2, \lambda_3, \lambda_4}^{\lambda}, \overbrace{\mu_1, \mu_2, \mu_3, \mu_4}^{\mu} : \lambda_3 \geq 0 \wedge \lambda_4 \geq 0 \wedge \mu_3 \geq 0 \wedge \mu_4 \geq 0 \wedge$$

$$\begin{pmatrix} \lambda \\ \mu \end{pmatrix} \begin{pmatrix} t_{xx} & t_{xy} & -1 & 0 \\ t_{yx} & t_{yy} & 0 & -1 \\ 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} -1 & -1 & 1 & 0 \\ 1 & 1 & -1 & 0 \end{pmatrix} \wedge \begin{pmatrix} \lambda \\ \mu \end{pmatrix} \begin{pmatrix} -t_x \\ -t_y \\ -1 \\ -1 \end{pmatrix} = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$$

This constraint requires that the right-hand side on the implication is obtained as a linear combination of the (in)equalities on the left-hand side of the implication. We conjoin the above constraint with *Constraint*.

We apply an SMT solver to compute a satisfying valuation of template parameters occurring in *Constraint* and obtain:

$$\frac{t_{xx} \mid t_{xy} \mid t_x \mid t_{yx} \mid t_{yy} \mid t_y}{1 \mid 1 \mid 0 \mid 0 \mid 0 \mid 10}$$

By applying *CexSol* on the template $v' = Tv + t$, which is the result of $\text{RELT}(\text{rel})(v, v')$, we obtain the conjunction $x' = x + y \wedge y' = 10$. In this example, we assume that the template $\text{GRDT}(\text{grd})(v)$ is equal to *true*. Hence, we modify the clauses that record the current candidate for $\text{rel}(v, v')$ and $\text{grd}(v)$ as follows:

$$\text{Defs} = \{x' = x + y \wedge y' = 10 \rightarrow \text{rel}(v, v'), \text{grd}(v) \rightarrow \text{true}\}$$

Now we proceed with the next iteration of the main loop in E-HSF.

Second candidate for Skolem relation The second iteration in E-HSF uses *Defs* and *Constraint* as determined during the first iteration. We apply HSF on $Skolemized \cup Defs$ and obtain an error derivation *Cex* consisting of the clauses

$$\begin{aligned} &init(v) \wedge q_1(v), \quad q_1(v) \wedge \neg dst(v) \wedge q_2(v, v') \rightarrow q_3(v, v'), \\ &x' = x + y \wedge y' = 10 \rightarrow q_2(v, v'), \quad q_3(v, v') \rightarrow q_4(v, v'), \end{aligned}$$

together with the function SYM such that $SYM(q_1) = inv$, $SYM(q_2) = rel$, $SYM(q_3) = rank$, and $SYM(q_4) = ti$. From *Cex* we extract $CexDefs = \{x' = x + y \wedge y' = 10 \rightarrow q_2(v, v')\}$ since $SYM(q_2) \in Rels$. We apply **RESOLVE** on $Cex \setminus CexDefs$ and obtain:

$$init(v) \wedge \neg dst(v) \wedge q_2(v, v') \rightarrow q_4(v, v') .$$

As seen at the first iteration, we have $RELT(rel)(v, v') = (v' = Tv + t)$. Hence we have $body = (init(v) \wedge \neg dst(v) \wedge v' = Tv + t)$.

Since $SYM(q_4) = ti$ and $dwf(ti) \in Skolemized$, the error derivation witnesses a violation of disjunctive well-foundedness. Hence, by applying **BOUND**T and **DECREASE**T we construct templates $bound(v)$ and $decrease(v, v')$ corresponding to a bound and decrease condition over the program variables, respectively.

$$\begin{aligned} bound(v) &= (r_x x + r_y y \geq r_0) , \\ decrease(v, v') &= (r_x x' + r_y y' \leq r_x x + r_y y - 1) . \end{aligned}$$

Finally, we set *head* to the conjunction $r_x x + r_y y \geq r_0 \wedge r_x x' + r_y y' \leq r_x x + r_y y - 1$.

By **ENCODE**VALIDITY on the implication $body \rightarrow head$ we obtain the constraint

$$\begin{aligned} &\exists \overbrace{\lambda_1, \lambda_2, \lambda_3, \lambda_4}^{\lambda}, \overbrace{\mu_1, \mu_2, \mu_3, \mu_4}^{\mu} : \lambda_3 \geq 0 \wedge \lambda_4 \geq 0 \wedge \mu_3 \geq 0 \wedge \mu_4 \geq 0 \wedge \\ &\quad \begin{pmatrix} \lambda \\ \mu \end{pmatrix} \begin{pmatrix} t_{xx} & t_{xy} & -1 & 0 \\ t_{yx} & t_{yy} & 0 & -1 \\ 0 & -1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{pmatrix} = \begin{pmatrix} -r_x & -r_y & 0 & 0 \\ -r_x & -r_y & r_x & r_y \end{pmatrix} \wedge \begin{pmatrix} \lambda \\ \mu \end{pmatrix} \begin{pmatrix} -t_x \\ -t_y \\ -1 \\ -1 \end{pmatrix} = \begin{pmatrix} -r_0 \\ -1 \end{pmatrix} . \end{aligned}$$

We add the above constraint as an additional conjunct to *Constraint*. That is, *Constraint* is strengthened during each iteration.

We apply the SMT solver to compute a valuation template parameters that satisfies *Constraint*. We obtain the following solution *CexSol*:

$$\begin{array}{c|c|c|c|c|c} t_{xx} & t_{xy} & t_x & t_{yx} & t_{yy} & t_y \\ \hline 1 & 0 & 1 & 0 & 0 & 1 \end{array}$$

The corresponding values of r and r_0 are $(-1, 0)$ and -1 , which lead to the bound $-x \geq 1$ and the decrease relation $-x' \leq -x - 1$. By applying *CexSol* on the template $v' = Tv + t$ we obtain the conjunction $x' = x + 1 \wedge y' = 1$. Note that the solution for $rel(v, v')$ obtained at this iteration is not compatible with

the solution obtained at the first iteration, i.e., the intersection of the respective Skolem relations is empty. Finally, we modify $Defs$ according to $CexSol$ and obtain:

$$Defs = \{x' = x + 1 \wedge y' = 1 \rightarrow rel(v, v'), \quad grd \rightarrow true\}$$

Now we proceed with the next iteration of the main loop in E-HSF. At this iteration the application of HSF returns a solution $ClauseSol$ such that

$$\begin{aligned} ClauseSol(inv(v)) &= (y \geq 1) , \\ ClauseSol(rel(v)) &= (x' = x + 1 \wedge y' = 1) , \\ ClauseSol(rank(v, v')) &= (x \leq -1 \wedge x' \geq x + 1) , \\ ClauseSol(ti(v, v')) &= (x \leq -1 \wedge x' \geq x + 1) . \end{aligned}$$

Thus, the algorithm E-HSF finds a solution to the original set of forall-exists Horn clauses (and hence proves the program satisfies the CTL property).

5 Verifying CTL properties using E-HSF

In this section we show how E-HSF can be used for automatically proving CTL properties of programs. We utilize a standard reduction step from CTL properties to existentially quantified Horn-like clauses with well-foundedness conditions, see e.g. [25]. Here, due to space constraints, we only illustrate the reduction, using examples and refer to [25] for details of the CTL proof system.

We consider a program over variables v , with an initial condition given by an assertion $init(v)$, and a transition relation $next(v, v')$. Given a CTL property, we generate Horn-like clauses such that the property is satisfied if and only if the set of clauses is satisfiable.

The generation proceeds in two steps. The first step decomposes the property into sub-properties by following the nesting structure of the path quantifiers that occur in the property. As a result we obtain a set of simple CTL formulas that contain only one path quantifier. Each property is accompanied by a predicate that represents a set of program states that needs to be discovered.

As an example, we present the decomposition of $(init(v), next(v, v')) \models_{CTL} AG(EF(dst(v)))$, where $dst(v)$ is a first-order assertion over v . Since $EF(dst(v))$ is a sub-formula with a path quantifier as the outmost symbol, we introduce a fresh predicate $p(v)$ that is used to replace $EF(dst(v))$. Furthermore, we require that every computation that starts in a state described by $p(v)$ satisfies $EF(dst(v))$. Since the resulting CTL formulas do not have any nested path quantifiers we stop the decomposition process. The original verification question is equivalent to the existence of $p(v)$ such that $(init(v), next(v, v')) \models_{CTL} AG(p(v))$ and $(p(v), next(v, v')) \models_{CTL} EF(dst(v))$.

At the second step we consider each of the verification sub-questions obtained by decomposing the property and generate Horn-like clauses that constrain auxiliary sets and relations over program states. For $(init(v), next(v, v')) \models_{CTL} AG(p(v))$ we obtain the following clauses over an auxiliary predicate $inv_1(v)$:

$$init(v) \rightarrow inv_1(v), \quad inv_1(v) \wedge next(v, v') \rightarrow inv_1(v'), \quad inv_1(v) \rightarrow p(v).$$

Due to the existential path quantifier in $(p(v), next(v, v')) \models_{CTL} EF(dst(v))$ we obtain clauses that contain existential quantification. We deal with the eventuality by imposing a well-foundedness condition. The resulting clauses over auxiliary $inv_2(v)$, $rank(v, v')$, and $ti(v, v')$ are below (note that $dst(v)$ is a constraint, and hence can occur under negation).

$$\begin{aligned} p(v) &\rightarrow inv_2(v), & inv_2(v) \wedge \neg dst(v) &\rightarrow \exists v' : next(v, v') \wedge rank(v, v'), \\ rank(v, v') &\rightarrow ti(v, v'), & ti(v, v') \wedge rank(v, v') &\rightarrow ti(v, v''), \quad dwf(ti). \end{aligned}$$

Finally, the above clauses have a solution for $inv_1(v)$, $p(v)$, $inv_2(v)$, $rank(v, v')$, and $ti(v, v')$ if and only if $(init(v), next(v, v')) \models_{CTL} AG(EF(dst(v)))$. Then, we apply E-HSF as a solver.

6 Experiments

In this section we present our implementation of E-HSF and its experimental evaluation on proving universal and existential CTL properties of programs.

Our implementation relies on HSF [19] to solve universally-quantified Horn clauses over linear inequalities (see line 4 in Figure 2) and on the Z3 solver [15] at line 19 in Figure 2 to solve (possibly non-linear) constraints. The input to our tool is a transition system described using Prolog facts $init(v)$ and $next(v, v')$, as well as forall-exists Horn clauses corresponding to the CTL property to be proved or disproved.

We run E-HSF on the examples from industrial code from [11, Figure 7]: **OS frag.1**, **OS frag.2**, **OS frag.3**, **OS frag.4**, **OS frag.5**, **PgSQL arch** and **S/W Updates**. For each pair of a program and CTL property ϕ , we generated two verification tasks: prove ϕ and prove $\neg\phi$. The existence of a proof for a property ϕ implies that $\neg\phi$ is violated by the same program. (Similarly, a proof for $\neg\phi$ implies that ϕ is violated by the same program.)

GRDT and RELT are provided by the user and need to satisfy Equation 1. Currently, this condition is not checked by the implementation, but could be done for linear templates using quantifier elimination techniques. For our examples, linear templates are sufficiently expressive. We use $RELT(next)(v, v') = (next(v, v') \wedge w' = Tv + t \wedge Gv \leq g)$ and $GRDT(next)(v, v') = (Gv \leq g \wedge \exists v' : next(v, v'))$, where w' is a subset of v that is left unconstrained by $next(v, v')$. Such w' are explicitly marked in the original benchmark programs using names **rho1**, **rho2**, \dots . For direct comparison with the results from [11], we used template functions corresponding to the **rho**-variables. The quantifier elimination in $\exists v' : next(v, v')$ can be automated for the theory of linear arithmetic. For dealing with well-foundedness we use linear ranking functions, and hence corresponding linear templates for **DECREASET** and **BOUNDT**.

We report the results in Table 1. Columns 3 and 6 show \checkmark marks for the cases where E-HSF was able to find a solution, i.e., prove the CTL property. See Columns 4 and 7 for the time spent on finding solutions. E-HSF is able to find proofs for all the correct programs except for P14 and P15 that correspond to

Program	Property ϕ	$\models_{CTL} \phi$			$\models_{CTL} \neg\phi$		
		Result	Time	Name	Result	Time	Name
P1	$AG(a = 1 \rightarrow AF(r = 1))$	✓	1.2s	1	×	2.7s	29
P2	$EF(a = 1 \wedge EG(r \neq 5))$	✓	0.6s	30	×	5.2s	2
P3	$AG(a = 1 \rightarrow EF(r = 1))$	✓	4.8s	3	×	0.1s	31
P4	$EF(a = 1 \wedge AG(r \neq 1))$	✓	0.6s	32	×	0.4s	4
P5	$AG(s = 1 \rightarrow AF(u = 1))$	✓	6.1s	5	×	0.2s	33
P6	$EF(s = 1 \wedge EG(u \neq 1))$	✓	1.4s	34	×	3.6s	6
P7	$AG(s = 1 \rightarrow EF(u = 1))$	✓	12.9s	7	×	0.2s	35
P8	$EF(s = 1 \wedge AG(u \neq 1))$	✓	44.7s	36	×	3.8s	8
P9	$AG(a = 1 \rightarrow AF(r = 1))$	✓	51.3s	9	×	120.0s	37
P10	$EF(a = 1 \wedge EG(r \neq 1))$	✓	132.0s	38	×	45.9s	10
P11	$AG(a = 1 \rightarrow EF(r = 1))$	✓	67.6s	11	×	3.9s	39
P12	$EF(a = 1 \wedge AG(r \neq 1))$	✓	67.9s	12	×	3.8s	40
P13	$AF(io = 1) \vee AF(ret = 1)$	✓	37m54s	13	T/O	-	41
P14	$EG(io \neq 1) \wedge EG(ret \neq 1)$	T/O	-	42	×	136.6s	14
P15	$EF(io = 1) \wedge EF(ret = 1)$	T/O	-	15	×	1.4s	43
P16	$AG(io \neq 1) \vee AG(ret \neq 1)$	✓	0.1s	44	×	874.5s	16
P17	$AG(AF(w \geq 1))$	✓	3.0s	17	×	0.1s	45
P18	$EF(EG(w < 1))$	✓	0.5s	46	×	3.5s	18
P19	$AG(EF(w \geq 1))$	✓	3.3s	19	×	0.1s	47
P20	$EF(AG(w < 1))$	✓	0.7s	48	×	0.1s	20
P21	$AG(AF(w = 1))$	✓	2.8s	21	×	0.1s	49
P22	$EF(EG(w \neq 1))$	✓	2.2s	50	×	5.0s	22
P23	$AG(EF(w = 1))$	✓	4.5s	23	×	0.1s	51
P24	$EF(AG(w \neq 1))$	✓	3.4s	52	×	0.7s	24
P25	$c > 5 \rightarrow AF(r > 5)$	✓	3.2s	25	×	0.1s	53
P26	$c > 5 \wedge EG(r \leq 5)$	×	0.1s	54	×	1.3s	26
P27	$c > 5 \rightarrow EF(r > 5)$	×	0.2s	27	×	0.1s	55
P28	$c > 5 \wedge AG(r \leq 5)$	×	0.1s	56	×	0.3s	28

Table 1. Evaluation of E-HSF on industrial benchmarks from [11]. Each “Name” column gives the corresponding program number in [11, Figure 7]. For P12, E-HSF returns different results compared to [11]. For P26, P27 and P28 both properties ϕ and $\neg\phi$ are satisfied only for some initial states. (Neither ϕ nor $\neg\phi$ hold for these programs.)

WINDOWS FRAG.4. Currently, E-HSF models the control flow symbolically using a program counter variable, which is most likely the reason for not succeeding on P14 and P15. Efficient treatment of control flow along the lines of explicit analysis as performed in the CPAchecker framework could lead to significant improvements for dealing with programs with large control-flow graphs [4].

For cases where the property contains more than one path quantifier and the top-most temporal quantifier is F or U , our implementation generates non-Horn clauses following the proof system from [25]. While a general algorithm for solving non-Horn clauses is beyond the scope of this paper, we used a simple heuristic to seed solutions for queries appearing under the negation operator. For example, for the verification task obtained from proving ϕ for P2, we used the solution $(a = 1 \wedge r \neq 5)$ for the query corresponding to the nesting structure of ϕ . This solution is obtained as a conjunction of the atomic constraints from ϕ .

7 Related work

Our work is inspired by a recent approach to CTL verification of programs [11]. The main similarity lies in the use of a refinement loop to discover witnesses for resolving non-determinism/existentially quantified variables. The main difference lies in the way candidate witnesses are selected. While [11] refines witnesses, i.e., the non-determinism in witness relations monotonically decreases at each iteration, E-HSF can change witness candidates arbitrarily (yet, subject to the global constraint). Thus, our method can backtrack from wrong choices in cases when [11] needs to give up.

E-HSF generalizes solving methods for universally quantified Horn clauses over decidable theories, e.g. [19, 23, 27]. Our approach relies on the templates for describing the space of candidate witnesses. Computing witnesses using a generalisation approach akin to PDR [23] is an interesting alternative to explore in future work.

Template based synthesis of invariants and ranking functions is a prominent technique for dealing with universal properties, see e.g. [10, 21, 30, 33]. E-HSF implementation of ENCODEVALIDITY supporting linear arithmetic inequalities is directly inspired by these techniques, and puts them to work for existential properties.

Decision procedures for quantified propositional formulas on bit as well as word level [24, 34] rely on iteration and refinement for the discovery of witnesses. The possibility of integration of QBF solvers as an implementation of ENCODEVALIDITY is an interesting avenue for future research.

Some formulations of proof systems for mu-calculus, e.g., [14] and [28], could be seen as another source of forall-exists clauses (to pass to E-HSF). Compared to the XSB system [14] that focuses on finite state systems, E-HSF aims at infinite state systems and employs a CEGAR-based algorithm. XSB's extensions for infinite state systems are rather specific, e.g., data-independent systems, and do not employ abstraction refinement techniques. Finally, we remark that abstraction-based methods, like ours, can be complemented with program specialization-based methods for verification of CTL properties [17].

8 Conclusion

Verification conditions for proving existential temporal properties of programs can be represented using existentially quantified Horn-like clauses. In this paper we presented a counterexample guided method for solving such clauses, which can compute witnesses to existentially quantified variables in form of linear arithmetic expressions. By aggregating constraints on witness relations across different counterexamples our method can recover from wrong choices. We leave the evaluation of applicability of our method for other problems requiring witness computation, e.g., software synthesis or game solving to future work.

Acknowledgements We thank Byron Cook and Eric Koskinen for valuable discussion and for generously making their benchmarks available. This research was supported in part by ERC project 308125 VeriSynth and by the DFG Graduiertenkolleg 1480 (PUMA).

References

1. A. Albarghouthi, Y. Li, A. Gurfinkel, and M. Chechik. Ufo: A framework for abstraction- and interpolation-based software verification. In *CAV*, 2012.
2. T. Ball and S. K. Rajamani. The SLAM project: debugging system software via static analysis. In *POPL*, 2002.
3. D. Beyer and M. E. Keremoglu. CPAchecker: A tool for configurable software verification. In *CAV*, 2011.
4. D. Beyer and S. Löwe. Explicit-state software model checking based on CEGAR and interpolation. In *FASE*, 2013.
5. B. Blanchet, P. Cousot, R. Cousot, J. Feret, L. Mauborgne, A. Miné, D. Monniaux, and X. Rival. A static analyzer for large safety-critical software. In *PLDI*, 2003.
6. R. Bruttomesso, A. Cimatti, A. Franzén, A. Griggio, and R. Sebastiani. The MathSAT 4SMT solver. In *CAV*, 2008.
7. C. Cadar, D. Dunbar, and D. R. Engler. KLEE: Unassisted and automatic generation of high-coverage tests for complex systems programs. In *OSDI*, 2008.
8. E. M. Clarke, D. Kroening, and F. Lerda. A tool for checking ANSI-C programs. In *TACAS*, 2004.
9. E. M. Clarke, D. Kroening, N. Sharygina, and K. Yorav. SATABS: SAT-based predicate abstraction for ANSI-C. In *TACAS*, 2005.
10. M. Colón, S. Sankaranarayanan, and H. Sipma. Linear invariant generation using non-linear constraint solving. In *CAV*, 2003.
11. B. Cook and E. Koskinen. Reasoning about nondeterminism in programs. In *PLDI*, 2013.
12. B. Cook, A. Podelski, and A. Rybalchenko. Termination proofs for systems code. In *PLDI*, 2006.

13. P. Cousot and R. Cousot. Abstract interpretation: A unified lattice model for static analysis of programs by construction or approximation of fixpoints. In *POPL*, 1977.
14. B. Cui, Y. Dong, X. Du, K. N. Kumar, C. R. Ramakrishnan, I. V. Ramakrishnan, A. Roychoudhury, S. A. Smolka, and D. S. Warren. Logic programming and model checking. In *PLILP/ALP*, 1998.
15. L. M. de Moura and N. Bjørner. Z3: An efficient SMT solver. In *TACAS*, 2008.
16. F. Emmes, T. Enger, and J. Giesl. Proving non-looping non-termination automatically. In *IJCAR*, 2012.
17. F. Fioravanti, A. Pettorossi, M. Proietti, and V. Senni. Generalization strategies for the verification of infinite state systems. *Theory and Practice of Logic Programming*, 13:175–199, 2 2013.
18. S. Graf and H. Saïdi. Construction of abstract state graphs with PVS. In *CAV*, 1997.
19. S. Grebenschikov, N. P. Lopes, C. Popeea, and A. Rybalchenko. Synthesizing software verifiers from proof rules. In *PLDI*, 2012.
20. A. Gupta, T. A. Henzinger, R. Majumdar, A. Rybalchenko, and R.-G. Xu. Proving non-termination. In *POPL*, 2008.
21. A. Gupta, R. Majumdar, and A. Rybalchenko. From tests to proofs. In *TACAS*, 2009.
22. T. A. Henzinger, R. Jhala, R. Majumdar, and K. L. McMillan. Abstractions from proofs. In *POPL*, 2004.
23. K. Hoder, N. Bjørner, and L. de Moura. μZ - an efficient engine for fixed points with constraints. In *CAV*, 2011.
24. M. Janota, W. Klieber, J. Marques-Silva, and E. M. Clarke. Solving QBF with counterexample guided refinement. In *SAT*, 2012.
25. Y. Kesten and A. Pnueli. A compositional approach to CTL* verification. *Theor. Comput. Sci.*, 331(2-3):397–428, 2005.
26. K. L. McMillan. Lazy abstraction with interpolants. In *CAV*, 2006.
27. K. L. McMillan and A. Rybalchenko. Computing relational fixed points using interpolation. Technical report, 2012. available from authors.
28. K. S. Namjoshi. Certifying model checkers. In *CAV*, 2001.
29. É. Payet and F. Spoto. Experiments with non-termination analysis for Java Bytecode. *Electr. Notes Theor. Comput. Sci.*, 253(5), 2009.
30. A. Podelski and A. Rybalchenko. A complete method for the synthesis of linear ranking functions. In *VMCAI*, 2004.
31. A. Podelski and A. Rybalchenko. Transition invariants. In *LICS*, 2004.
32. A. Schrijver. *Theory of linear and integer programming*. Wiley-Interscience series in discrete mathematics and optimization. Wiley, 1999.
33. S. Srivastava and S. Gulwani. Program verification using templates over predicate abstraction. In *PLDI*, 2009.
34. C. M. Wintersteiger, Y. Hamadi, and L. M. de Moura. Efficiently solving quantified bit-vector formulas. *Formal Methods in System Design*, 2013.