



# Sound and Complete Witnesses for Template-Based Verification of LTL Properties on Polynomial Programs

Krishnendu Chatterjee<sup>1</sup>, Amir Goharshady<sup>2</sup>, Ehsan Goharshady<sup>1</sup>,  
Mehrdad Karrabi<sup>1</sup>(✉), and Đorđe Žikelić<sup>3</sup>



<sup>1</sup> Institute of Science and Technology Austria (ISTA),  
Klosterneuburg, Austria

{krishnendu.chatterjee,ehsan.goharshady,  
mehrdad.karrabi}@ist.ac.at

<sup>2</sup> The Hong Kong University of Science and Technology (HKUST),  
Clear Water Bay, Hong Kong

goharshady@cse.ust.hk

<sup>3</sup> Singapore Management University, Singapore, Singapore  
dzikelic@smu.edu.sg

**Abstract.** We study the classical problem of verifying programs with respect to formal specifications given in the linear temporal logic (LTL). We first present novel sound and complete witnesses for LTL verification over imperative programs. Our witnesses are applicable to both verification (proving) and refutation (finding bugs) settings. We then consider LTL formulas in which atomic propositions can be polynomial constraints and turn our focus to polynomial arithmetic programs, i.e. programs in which every assignment and guard consists only of polynomial expressions. For this setting, we provide an efficient algorithm to automatically synthesize such LTL witnesses. Our synthesis procedure is both sound and semi-complete. Finally, we present experimental results demonstrating the effectiveness of our approach and that it can handle programs which were beyond the reach of previous state-of-the-art tools.

## 1 Introduction

*Linear-Time Temporal Logic.* The Linear-time Temporal Logic (LTL) [53] is one of the most classical and well-studied frameworks for formal specification, model checking and program verification. In LTL, we consider a set  $\text{AP}$  of atomic propositions and an infinite trace which tells us which propositions in  $\text{AP}$  hold at any given time. LTL formulas are then able to not only express propositional logical operations, but also modalities referring to the future. For example,  $X p$  requires that  $p$  holds in the next timeslot, whereas  $F q$  means  $q$  should hold at

---

Đorđe Žikelić: Part of the work done while the author was at the Institute of Science and Technology Austria (ISTA).

some time in the future. This allows LTL to express common verification tasks such as termination, liveness, fairness and safety.

**Witnesses.** Given a specification  $\varphi$  and a program  $P$ , a *witness* is a mathematical object whose existence proves that the specification  $\varphi$  is satisfied by  $P$ . We say that a witness family is *sound and complete* when for every program  $P$  and specification  $\varphi$ , we have  $P \models \varphi$  if and only if there is a witness in the family that certifies it. Witnesses are especially useful in dealing with undecidable problems in verification, which includes all non-trivial semantic properties [56]. This is because although the general case of the problem is undecidable, having a sound and complete notion of a witness can lead to algorithms that check for the existence of witnesses of a special form. For example, while termination is undecidable [64], and hence so is the equivalent problem of deciding the existence of a ranking function, there are nevertheless sound and complete algorithms for synthesis of *linear* ranking functions [54]. Similarly, while reachability (safety violation) is undecidable, it has sound and complete witnesses that can be automatically synthesized in linear and polynomial forms [1]. Our work subsumes both [54] and [1] and provides sound and complete witnesses for general LTL formulas.

**Polynomial Programs.** In this work, we mainly focus on imperative programs with polynomial arithmetic. More specifically, our programs have real variables and the right-hand-side of every assignment is a polynomial expression with respect to program variables. Similarly, the guard of every loop or branch is also a boolean combination of polynomial inequalities over the program variables.

**Our Contributions.** In this work, our contributions are threefold:

- On the theoretical side, by exploiting the connections to Büchi automata, we propose a novel family of sound and complete witnesses for general LTL formulas. This extends and unifies the known concepts of ranking functions [36], inductive reachability witnesses [1] and inductive invariants [25], which are sound and complete witnesses for termination, reachability and safety, respectively. Our theoretical result is not limited to polynomial programs.
- On the algorithmic side, we consider polynomial programs and present a sound and semi-complete template-based algorithm to synthesize polynomial LTL witnesses. This algorithm is a generalization of the template-based approaches in [1, 25, 54] which considered termination, reachability and safety. To the best of our knowledge, this is the most general model checking problem over polynomial programs to be handled by template-based approaches to date.
- Finally, on the experimental side, we provide an implementation of our approach and comparisons with state-of-the-art LTL model checking tools. Our experiments show that our approach is applicable in practice and can handle many instances that were beyond the reach of previous methods. Thus, our completeness result pays off in practice and enables us to solve new instances.

**Motivation for Polynomial Programs.** There are several reasons why we consider polynomial programs:

- Many real-world families of programs, such as, programs for cyber-physical systems and smart contracts, can be modeled in this framework [10, 38, 42].
- They are one of the most general families for which finding polynomial witnesses for reachability and safety are known to be decidable [1, 12, 57]. Hence, they provide a desirable tradeoff between decidability and generality.
- Using abstract interpretation, non-polynomial behavior in a program can be removed or replaced by non-determinism. Moreover, one can approximate any continuous function up to any desired level of accuracy by a polynomial. This is due to the Stone–Weierstrass theorem [30]. Thus, analysis of polynomial programs can potentially be applied to many non-polynomial programs via abstract interpretation or numerical approximation of the program’s behavior.
- Previous works have studied (a) linear/affine programs with termination, safety, and reachability specifications [25, 54, 58], and (b) polynomial programs with termination, safety and reachability properties [1, 11, 12, 57]. Since LTL subsumes all these specifications, polynomial program analysis with LTL provides a unifying and general framework for all these previous works.

**Related Works on Linear Programs.** There are many approaches focusing on linear witness synthesis for important special cases of LTL formulas. For example, [43, 54] consider the problem of synthesizing linear ranking functions (termination witnesses) over linear arithmetic programs. The works [25, 58] synthesize linear inductive invariants (safety witnesses), while [39] considers probabilistic reachability witnesses. The work [41] handles a larger set of verification tasks and richer settings, such as context-sensitive interprocedural program analysis. All these works rely on the well-known Farkas lemma [32] and can handle programs with linear/affine arithmetic and synthesize linear/affine witnesses. In comparison, our approach is (i) applicable to general LTL formulas and not limited to a specific formula such as termination or safety, and (ii) able to synthesize *polynomial* witnesses for *polynomial* programs with soundness and completeness guarantees. Thus, our setting is more general in terms of (a) formulas, (b) witnesses, and (c) programs that can be supported.

**Related Works on Polynomial Programs.** Similar to the linear case, there is a rich literature on synthesis of polynomial witnesses over polynomial programs. However, these works again focus on specific special formulas only and are not applicable to general LTL. For example, [11, 15, 16, 44, 49, 51, 59, 68] consider termination analysis, [12] extends the invariant generation (safety witness synthesis) algorithm of [25] to the polynomial case and [14, 17, 18, 35, 62, 69] add support for probabilistic programs. The works [22, 70, 71] consider alternative types of witnesses for safety (barriers) and obtain similarly successful synthesis algorithms. Finally, [1, 63] synthesize reachability witnesses. Since we can handle any arbitrary LTL formula, our approach can be seen as an extension and unification of all these works. Indeed, our synthesis algorithm directly builds upon and extends [1].

In both cases above, some of the previous works are incomparable to ours since they consider probabilistic programs, whereas our setting has only non-probabilistic polynomial programs. Note that we do allow non-determinism.

**Related Works on LTL Model Checking.** There are thousands of works on LTL model checking and there is no way we can do justice to all. We refer to [24, 60] for an excellent treatment of the finite-state cases. Some works that provide LTL model checking over infinite-state systems/programs are as follows:

- A prominent technique in this area is predicate abstraction [29, 40, 55], which uses a finite set of abstract states defined by an equivalence relation based on a finite set of predicates to soundly, but not completely, reduce the problem to the finite-state case.
- [19] uses a compositional approach to falsify LTL formulas and find an indirect description of a path that violates the specification.
- There are several symbolic approaches, including [26] which is focused on fairness and [4] which is applicable to LLVM. Another work in this category is [31], whose approach is to repeatedly rule out infeasible finite prefixes in order to find a run of the program that satisfies/violates the desired LTL formula. The work [27] uses CTL-based approaches that might report false counter-examples when applied to LTL. It then identifies and removes such spurious counterexamples using symbolic determinization.
- The work [33] presents a framework for proving liveness properties in multi-threaded programs by using well-founded proof spaces.
- The recent work [52] uses temporal prophecies, inspired by classical prophecy variables, to provide significantly more precise reductions from general temporal verification to the special case of safety.
- There are many tools for LTL-based program analysis. For example, T2 [8] is able to verify a large family of liveness and safety properties, nuXmv [20] is a symbolic model checker with support for LTL, F3 [19] proves fairness in infinite-state transition systems, and Ultimate LTLAutomizer [31] is a general-purpose tool for verification of LTL specifications over a wide family of programs with support for various types of variables.
- Finally, we compare against the most recent related work [65]. This work provides relative-completeness guarantees for general programs with LTL specifications. Since it considers integer programs with recursive functions, there is no complexity guarantee provided. The earlier work [66] provides several special cases where termination is guaranteed. However, no runtime bounds are established. In contrast, our approach has both termination guarantees and sub-exponential time complexity for fixed degree.

As shown by our experimental results in Sect. 5, our completeness results enable our tool to handle instances that other approaches could not. On the other hand, our method is limited to polynomial programs and witnesses. Thus, there are also cases in which our approach fails but some of the previous tools succeed, e.g. when the underlying program requires a non-polynomial witness. In particular, Ultimate LTLAutomizer [31] is able to handle non-polynomial programs and witnesses, too.

## 2 Transition Systems, LTL and Büchi Automata

For a vector  $e \in \mathbb{R}^n$ , we use  $e_i$  to denote the  $i$ -th component of  $e$ . Given a finite set  $\mathcal{V}$  of real-valued variables, a variable valuation  $e \in \mathbb{R}^{|\mathcal{V}|}$  and a boolean predicate  $\varphi$  over  $\mathcal{V}$ , we write  $e \models \varphi$  when  $\varphi$  evaluates to true upon substituting variables by the values given in  $e$ .

We consider imperative numerical programs with real-valued variables, containing standard programming constructs such as assignments, branching and loops. In addition, our programs can have finite non-determinism. We denote non-deterministic branching in our syntax by **if \* then**. See Fig. 1 for an example. We use transition systems to formally model programs.

**Transition Systems.** An infinite-state *transition system* is a tuple  $\mathcal{T} = (\mathcal{V}, L, l_{init}, \theta_{init}, \mapsto)$ , where:

- $\mathcal{V} = \{x_0, \dots, x_{n-1}\}$  is a finite set of real-valued *program variables*.
- $L$  is a finite set of *locations* with  $l_{init} \in L$  the *initial location*.
- $\theta_{init} \subseteq \mathbb{R}^n$  is a set of *initial variable valuations*.
- $\mapsto$  is a finite set of *transitions*. Each transition  $\tau \in \mapsto$  is of the form  $\tau = (l, l', G_\tau, U_\tau)$ , where  $l$  is the source location,  $l'$  is the target location,  $G_\tau$  is the guard of the transition, which is a boolean predicate over  $\mathcal{V}$ , and  $U_\tau : \mathbb{R}^n \rightarrow \mathbb{R}^n$  is the update function of the transition.

Translating programs into transition systems is a standard process. In what follows, we assume we are given a transition system  $\mathcal{T} = (\mathcal{V}, L, l_{init}, \theta_{init}, \mapsto)$  of the program that we wish to analyze. An example is shown in Fig. 1.

**States and Runs.** A *state* in  $\mathcal{T}$  is a pair  $(l, e)$  with  $l \in L$  and  $e \in \mathbb{R}^n$ . A state  $(l, e)$  is said to be *initial* if  $l = l_{init}$  and  $e \in \theta_{init}$ . We use  $\mathcal{S}$  and  $\mathcal{S}_{init}$  to denote the sets of all states and initial states. We assume the existence of a special *terminal location*  $l_t$  with a single outgoing transition which is a self-loop  $(l_t, l_t, \text{true}, Id)$  with  $Id(e) = e$  for each  $e \in \mathbb{R}^n$ . A state  $(l', e')$  is a *successor* of  $(l, e)$ , denoted as  $(l, e) \mapsto (l', e')$ , if there exists a transition  $\tau = (l, l', G_\tau, U_\tau) \in \mapsto$  such that  $e \models G_\tau$  and  $e' = U_\tau(e)$ . We assume each state has at least one successor so that all runs are infinite and LTL semantics are well defined. This is without loss of generality, since we can introduce transitions to the terminal location. A *run* in  $\mathcal{T}$  is an infinite sequence of successor states starting in  $\mathcal{S}_{init}$ .

**Linear-Time Temporal Logic (LTL).** Let AP be a finite set of atomic propositions. LTL formulas are inductively defined as follows:


- If  $p \in \text{AP}$ , then  $p$  is an LTL formula.
- If  $\varphi$  and  $\psi$  are LTL formulas, then  $\neg\varphi$ ,  $\varphi \vee \psi$ ,  $\varphi \wedge \psi$ , **X**  $\varphi$ , **G**  $\varphi$ , **F**  $\varphi$  and  $\varphi \text{ U } \psi$  are all LTL formulas.

$\neg, \vee$  and  $\wedge$  are the propositional negation, disjunction and conjunction while **X**, **G**, **F** and **U** are the *next*, *globally*, *finally* and *until* temporal operators.

**Atomic Propositions.** To use LTL over the transition system  $\mathcal{T}$ , we first need to specify a finite set of atomic propositions AP. In this work, we let the set AP

```

Precondition ( $l_{init}$ ):  $x_0 \geq 0$ 
 $l_1$ : while  $x_0 \geq 0$  do
    if * then
 $l_2$ :    $x_0 = x_0^2 + 1$ 
    else
 $l_3$ :    $x_0 = x_0^2 - 1$ 
 $l_t$ :
    
```



**Fig. 1.** An example program (left) and its transition system (right). Note that there is non-determinism at  $l_1$ .

consist of (i) finitely many constraints of the form  $\text{exp}(\mathbf{x}) \geq 0$  where  $\text{exp}: \mathcal{V} \rightarrow \mathbb{R}$  is an arithmetic expression over  $\mathcal{V}$ , and (ii) an atomic proposition  $at(l)$  for each location  $l$  in  $\mathcal{T}$ . Note that unlike classical LTL settings, our atomic propositions are not necessarily independent. For example, if we have  $p_1 := x \geq 0$  and  $p_2 := x + 1 \geq 0$ , it is impossible to have  $p_1 \wedge \neg p_2$  at any point in time.

The semantics of LTL is standard, refer to the extended version of the paper [13] for details.

**Program Analysis with LTL Specifications.** We now define the LTL program analysis problems that we consider in this work. Given a transition system  $\mathcal{T}$  and an LTL formula  $\varphi$ , we are interested in two problems:

1. *LTL Verification of Programs (LTL-VP).* Given a transition system  $\mathcal{T}$  and an LTL formula  $\varphi$  in  $\mathcal{T}$ , prove that *all possible runs* of  $\mathcal{T}$  satisfy  $\varphi$ .
2. *LTL Refutation of Programs (LTL-RP).* Given a transition system  $\mathcal{T}$  and an LTL formula  $\varphi$  in  $\mathcal{T}$ , prove that there *exists a run* that violates  $\varphi$ , or equivalently, satisfies  $\neg\varphi$ .

**Remark.** *LTL Verification* asks about correctness of the program while *LTL Refutation* addresses the problem of finding bugs. Both problems have been widely studied in the literature [3, 31, 65]. Moreover, a witness for the refutation problem can be used in counterexample-guided techniques such as CEGAR [23].

**Example.** Consider the transition system in Fig. 1 and the LTL formula  $\varphi = \neg[G(at(l_3) \Rightarrow Fat(l_2))]$ . The run that starts at  $(l_{init}, 1)$  and chooses  $l_2$  if  $x_0 = 0$  and  $l_3$  whenever  $x_0 = 1$ , does not satisfy  $\varphi$ . Therefore, in this case, the answer to the LTL-RP problem is positive. Additionally, deciding termination of a program with terminal location  $l_t$  is equivalent to the LTL-VP problem of  $[F at(l_t)]$  on the same program.

**Program Analysis with Büchi Specifications.** A Büchi specification is a subset  $\mathcal{B} \subseteq \mathcal{S}$  of states. A run  $\pi$  is  $\mathcal{B}$ -Büchi if it visits  $\mathcal{B}$  infinitely many times, i.e. if  $\{i \mid \pi(i) \in \mathcal{B}\}$  is infinite. Similar to LTL, Büchi specifications give rise to two main decision problems as follows:

1. *Universal Büchi Program Analysis (UB-PA).* Given a transition system  $\mathcal{T}$  and a Büchi specification  $\mathcal{B}$  on  $\mathcal{T}$ , prove that *all possible runs* of  $\mathcal{T}$  are  $\mathcal{B}$ -Büchi.


2. *Existential Büchi Program Analysis (EB-PA)*. Given a transition system  $\mathcal{T}$  and a Büchi specification  $\mathcal{B}$  on  $\mathcal{T}$ , prove the *existence* of a run that is  $\mathcal{B}$ -Büchi.

**Büchi Automata** [2,9]. A *non-deterministic Büchi automaton (NBW)* is a tuple  $N = (Q, A, \delta, q_0, F)$ , where  $Q$  is a finite set of states,  $A$  is a finite alphabet,  $\delta: Q \times A \rightarrow 2^Q$  is a transition relation,  $q_0$  is the initial state, and  $F \subseteq Q$  is the set of accepting states. An infinite word  $a_0, a_1, \dots$  of letters in the alphabet  $A$  is accepted by  $N$  if it gives rise to at least one accepting run in  $N$ , i.e. if there exists a run  $q_0, q_1, \dots$  such that  $q_{i+1} \in \delta(q_i, a_i)$  for each  $i$  and  $F$  is visited infinitely many times. It is a classical result that for every LTL formula  $\varphi$  defined over atomic predicates  $AP$  there exists a non-deterministic Büchi automaton  $N$  with alphabet  $2^{AP}$  which accepts exactly those traces that satisfy  $\varphi$  [24].

Let  $\mathcal{T} = (\mathcal{V}, L, l_{init}, \theta_{init}, \mapsto)$  be a transition system and  $N = (Q, 2^{AP}, \delta, q_0, F)$  be an NBW. In order to analyse  $\mathcal{T}$  with respect to  $N$ , we utilize the Cartesian product  $\mathcal{T} \times N$  and the Büchi specification  $\mathcal{B}_N^{\mathcal{T}} = L \times F \times \mathbb{R}^n$ . The state space of  $\mathcal{T} \times N$  is exactly the Cartesian product of the state spaces of  $\mathcal{T}$  and  $N$ . Moreover, for  $l, l' \in L$  and  $q, q' \in Q$ , there is a transition from  $(l, q)$  to  $(l', q')$  if there is a transition in  $\mathcal{T}$  from  $l$  to  $l'$  and a transition in  $N$  from  $q$  to  $q'$ . The formal definition of the product is available in [13]. See Fig. 2 for an example.

**Lemma 1 (From LTL to Büchi Specifications, Proof in [13]).** *Let  $\mathcal{T}$  be a transition system,  $\varphi$  an LTL formula for  $\mathcal{T}$  and  $N$  an NBW that accepts the same language as  $\varphi$ .*

- *The LTL-RP problem of  $\mathcal{T}$  and  $\neg\varphi$  is equivalent to the EB-PA problem of  $\mathcal{T} \times N$  and  $\mathcal{B}_N^{\mathcal{T}}$  [31].*
- *If  $N$  is deterministic, then the LTL-VP problem of  $\mathcal{T}$  and  $\varphi$  is equivalent to the UB-PA problem of  $\mathcal{T} \times N$  and  $\mathcal{B}_N^{\mathcal{T}}$ .*



**Fig. 2.** An NBW accepting  $GF\ at(l_2)$  with gray accepting nodes (left) and the product of the transition system in Fig. 1 and this NBW (right). A node labeled  $i, j$  represents location  $(l_i, q_j)$ . Unreachable locations have been removed. (Color figure online)

**Remark.** Based on the lemma above, instead of designing witnesses for the LTL-RP problem, we only need to find sound and complete witnesses for EB-PA. Moreover, it is easy to see that LTL-VP is reducible to LTL-RP since all

runs of  $\mathcal{T}$  satisfy  $\varphi$  if and only if there is no run that satisfies  $\neg\varphi$ . So, finding sound and complete witnesses for EB-PA will theoretically solve both verification and refutation variants of LTL program analysis. Note that the second statement in Lemma 1 is more restrictive than the first one since it only applies to deterministic Büchi automata. Thus, if the LTL formula  $\varphi$  does not admit a deterministic Büchi automaton, the above sequence of reductions from LTL-VP to LTL-RP should be made and then the EB-PA witness should be used. However, if  $\varphi$  admits a DBW, then the reduction to UB-PA is preferable in practice. We will provide witness concepts for both EB-PA and UB-PA problems in the next section.

### 3 Sound and Complete B-PA Witnesses

Let  $\mathcal{T} = (\mathcal{V}, L, l_{init}, \theta_{init}, \mapsto)$  be a transition system and  $\mathcal{B} \subseteq \mathcal{S}$  a set of states in  $\mathcal{T}$ . In this section, we introduce our sound and complete witnesses for the EB-PA and UB-PA problems.

#### 3.1 Sound and Complete Witnesses for Existential B-PA

Our witness concept for the EB-PA problem is a function that assigns a real value to each state in  $\mathcal{T}$ . The witness function is required to be non-negative in at least one initial state of  $\mathcal{T}$ , to preserve non-negativity in at least one successor state and to strictly decrease in value in at least one successor state whenever the current state is not contained in  $\mathcal{B}$  and the value of the witness function in the current state is non-negative. Hence, starting in an initial state in which the witness function is non-negative, one can always select a successor state in which the witness function is non-negative and furthermore ensure that  $\mathcal{B}$  is eventually reached due to the strict decrease condition, which will also be referred to as the *Büchi-ranking condition*. Intuitively, an EBRF is a function that overestimates the distance to  $\mathcal{B}$  and guarantees that  $\mathcal{B}$  is reached along at least one program run, at every program state in which the value of the EBRF is non-negative.

**Definition 1 (EBRF).** *Given two states  $s_1, s_2 \in \mathcal{S}$ , a function  $f: \mathcal{S} \rightarrow \mathbb{R}$  is said to Büchi-rank  $(s_1, s_2)$  where  $s_1 \mapsto s_2$ , if it satisfies one of the following:*

- $s_1 \in \mathcal{B} \wedge [f(s_1) \geq 0 \Rightarrow f(s_2) \geq 0]$ ; or
- $s_1 \notin \mathcal{B} \wedge [f(s_1) \geq 0 \Rightarrow 0 \leq f(s_2) \leq f(s_1) - 1]$ .

*$f$  is called a  $\mathcal{B}$ -Existential Büchi Ranking Function ( $\mathcal{B}$ -EBRF) if it satisfies the following conditions:*

- $\exists s_{init} \in \mathcal{S}_{init}$  where  $f(s_{init}) \geq 0$ .
- For every  $s_1 \in \mathcal{S}$ , there exists  $s_2 \in \mathcal{S}$  such that  $s_1 \mapsto s_2$  and  $(s_1, s_2)$  is Büchi-ranked by  $f$ .



**Example.** The following is a  $\{(l_1, q_1, *)\}$ -EBRF for the transition system in Fig. 2:  $f(l, x_0) = x_0 + 3$  if  $l = (l_{init}, q_0)$ ,  $f(l, x_0) = x_0 + 2$  if  $l = (l_1, q_0)$ ,  $f(l, x_0) = x_0 + 1$  if  $l = (l_2, q_0)$ ,  $f(l, x_0) = 0$  if  $l = (l_1, q_1)$  and  $f(l, x_0) = 0$  otherwise.

For example, the state  $s_0 = ((l_1, q_0), 1)$  has two successors in the transition system:  $s_1 = ((l_2, q_0), 1)$  and  $s_2 = ((l_3, q_0), 1)$ . It is easy to see that  $0 \leq f(s_1) \leq f(s_0) - 1$  which shows that transition from  $s_0$  to  $s_1$  is Büchi-ranked by  $f$ .

The following theorem, proved in the extended version [13], establishes the soundness and completeness of EBRFs for the EB-PA problem, which is the main result of this section. Hence, since we showed in Lemma 1 that one can reduce the LTL-RP problem to EB-PA, as a corollary it also follows that EBRFs provide sound and complete certificates for LTL-RP.

**Theorem 1 (Soundness and Completeness of EBRFs for EB-PA).**

*There exists a  $\mathcal{B}$ -EBRF  $f$  for  $\mathcal{T}$  with Büchi specification  $\mathcal{B}$  if and only if the answer to the EB-PA problem of  $\mathcal{T}$  and  $\mathcal{B}$  is positive.*

**Corollary 1.** *The answer to the LTL-RP problem of  $\mathcal{T}$  and  $\varphi$  is positive if and only if there exists a  $\mathcal{B}_N^T$ -EBRF for  $\mathcal{T} \times N$ , where  $N$  is the NBW accepting  $\neg\varphi$ .*

**3.2 Sound and Complete Witnesses for Universal B-PA**

Similarly to EBRFs, we can define a witness function for the UB-PA problem. The difference compared to EBRFs is that we now impose the Büchi ranking condition for *every* successor state of a state in which the witness function is non-negative. In contrast, in EBRFs we imposed the Büchi ranking condition only for *some* successor state.

**Definition 2 (UBRF).** *A function  $f: \mathcal{S} \rightarrow \mathbb{R}^n$  is called a  $\mathcal{B}$ -Universal Büchi Ranking Function ( $\mathcal{B}$ -UBRF) if it satisfies the following conditions:*

- $f(s) \geq 0$  for **every**  $s \in \mathcal{S}_{init}$
- For **every**  $s_1, s_2 \in \mathcal{S}$  such that  $s_1 \mapsto s_2$ ,  $(s_1, s_2)$  is Büchi-ranked by  $f$ .

We have the following theorem, which establishes that UBRFs provide a sound and complete certificate for the UB-PA problem. The proof is similar to the existential case and presented in the extended version [13]. The subsequent corollary then follows from Lemma 1 which shows that the LTL-VP problem can be reduced to the UB-PA problem if  $\varphi$  admits a deterministic Büchi automaton.

**Theorem 2 (Soundness and Completeness of UBRFs for UB-PA).**

*There exists a  $\mathcal{B}$ -UBRF  $f$  for  $\mathcal{T}$  with Büchi specification  $\mathcal{B}$  if and only if the answer to the UB-PA problem of  $\mathcal{T}$  and  $\mathcal{B}$  is positive.*

**Corollary 2.** *If  $\varphi$  is an LTL formula that admits a DBW  $D$ , the answer to the LTL-VP problem of  $\mathcal{T}$  and  $\varphi$  is positive iff there exists a  $\mathcal{B}_D^T$ -UBRF for  $\mathcal{T} \times D$ .*

**Remark.** Note that if the transition system  $\mathcal{T}$  is deterministic, (i.e. it contains no non-determinism in initial states, assignments or branches) the LTL-VP of  $\mathcal{T}$  and  $\varphi$  will be equivalent to the LTL-RP of  $\mathcal{T}$  and  $\neg\varphi$ . Thus, in this case, the Büchi automaton determinism assumption can be relaxed as follows: if  $N$  is an NBW that accepts the same language as  $\varphi$ , the answer to the LTL-VP of  $\mathcal{T}$  and  $\varphi$  is positive if and only if there exists a  $\mathcal{B}_N^T$ -EBRF for  $\mathcal{T} \times N$ .

## 4 Template-Based Synthesis of Polynomial Witnesses

We now present our fully automated algorithms to synthesize polynomial EBRFs and UBRFs in polynomial transition systems. A transition system  $\mathcal{T}$  is said to be *polynomial* if guards and updates of all transitions in  $\mathcal{T}$  are polynomial expressions over program variables  $\mathcal{V}$ . Given a polynomial transition system  $\mathcal{T}$  and a Büchi specification  $\mathcal{B}$ , which was obtained from an LTL formula as above, our approach synthesizes polynomial EBRFs and UBRFs of any desired degree, assuming that they exist. Our algorithms follow a template-based synthesis approach, similar to the methods used for reachability and termination analysis [1, 12]. In particular, both EBRF and UBRF synthesis algorithms first fix a symbolic polynomial template function for the witness at each location in  $\mathcal{T}$ . The defining conditions of EBRFs/UBRFs are then expressed as entailment constraint of the form

$$\exists c \in \mathbb{R}^m \quad \forall e \in \mathbb{R}^n \quad (\phi \Rightarrow \psi), \quad (1)$$

where  $\phi$  and  $\psi$  are conjunctions of polynomial inequalities. We show that this translation is sound and complete. However, such constraints are notoriously difficult to solve due to the existence of a quantifier alternation. Thus, we use the sound and semi-complete technique of [1] to eliminate the quantifier alternation and translate our constraints into a system of purely existentially quantified quadratic inequalities. Finally, this quadratic programming instance is solved by an SMT solver. We note that a central technical difficulty here is to come up with sound and complete witness notions whose synthesis can be reduced to solving entailment constraints of the form (1). While [1, 12] achieved this for termination and reachability, our EBRF and UBRF notions significantly extend these results to arbitrary LTL formulas.

As is common in static analysis tasks, we assume that the transition system comes with an invariant  $\theta_l$  at every location  $l$  in  $\mathcal{T}$ . Invariant generation is an orthogonal and well-studied problem. In polynomial programs, invariants can be automatically generated using the tools in [12, 34, 45]. Alternatively, one can encode an inductive invariant via constraints of the form (1). This has the extra benefit of ensuring that we always find an invariant that leads to a witness for our LTL formula, if such a witness exists, and thus do not sacrifice completeness due to potentially loose invariants. See [12] for details of the encoding. This is the route we took in our tool, i.e. our tool automatically generates the invariants it requires using the sound and complete method of [12]. For brevity, we removed the invariant generation part from the description of the algorithms below.

**Synthesis of Polynomial EBRFs.** We now present our algorithm for synthesis of a polynomial EBRF, given a polynomial transition system  $\mathcal{T} = (\mathcal{V}, L, l_{init}, \theta_{init}, \mapsto)$  and Büchi specification  $\mathcal{B}$  obtained from an LTL formula with polynomial inequalities in AP. We present a detailed example that illustrates the steps of the algorithm in the extended version of the paper [13]. The algorithm has five steps:

1. *Fixing Symbolic Templates.* Let  $M_{\mathcal{V}}^D = \{m_1, m_2, \dots, m_k\}$  be the set of all monomials of degree at most  $D$  over the set of variables  $\mathcal{V}$ . In the first step, the algorithm generates a symbolic polynomial template for the EBRF at each location  $l \in L$  as follows:  $f_l(x) = \sum_{i=1}^k c_{l,i} \cdot m_i$ . Here, all the  $c$ -variables are fresh symbolic template variables that represent the coefficients of polynomial expressions in  $f$ . The goal of our synthesis procedure is to find a concrete valuation of  $c$  variables for which  $f$  becomes a valid  $\mathcal{B}$ -EBRF for  $\mathcal{T}$ .
2. *Generating Entailment Constraints.* For every location  $l \in L$  and variable valuation  $x \models \theta_l$ , there must exist an outgoing transition  $\tau$  such that  $x \models G_\tau$  and  $\tau$  is Büchi-ranked by  $f$  in  $x$ . The algorithm symbolically writes down this condition as an entailment constraint:  $\forall x \in \mathbb{R}^n \quad x \models (\phi_l \Rightarrow \psi_l)$  with  $\phi_l$  and  $\psi_l$  symbolically computed as follows:  $\phi_l := \theta_l \wedge f_l(x) \geq 0$  and  $\psi_l \equiv \bigvee_{\tau \in \text{Out}_l} G_\tau \wedge \mathcal{B}\text{-Rank}(\tau)$ , where for each  $\tau = (l, l', G_\tau, U_\tau)$  the predicate  $\mathcal{B}\text{-Rank}$  is defined as

$$\mathcal{B}\text{-Rank} \equiv \begin{cases} f_{l'}(U_\tau(x)) \geq 0 \wedge f_{l'}(U_\tau(x)) \leq f_l(x) - 1 & l \notin \mathcal{B} \\ f_{l'}(U_\tau(x)) \geq 0 & l \in \mathcal{B} \end{cases}$$

The algorithm then writes  $\psi_l$  in disjunctive normal form as  $\bigvee_{i=1}^k \psi_{l,i}$ . Next, the algorithm rewrites  $\phi_l \Rightarrow \psi_l$  equivalently as:

$$(\phi_l \wedge \bigwedge_{i=1}^{k-1} \neg \psi_{l,i}) \Rightarrow \psi_{l,k} \quad (2)$$

This rewriting makes sure that we can later manipulate the constraint in (2) to fit in the standard form of (1)<sup>1</sup>. Intuitively, (2) ensures that whenever  $l$  was reached and each of the first  $k-1$  outgoing transitions were either unavailable or not Büchi-ranked by  $f$ , then the last transition has to be available and Büchi-ranked by  $f$ . Our algorithm populates a list of all constraints and adds the constraint (2) to this list before moving to the next location and repeating the same procedure. Note that in all of the generated constraints of the form (2), both the LHS and the RHS of the entailment are boolean combinations of polynomial inequalities over program variables.

3. *Reduce Constraints to Quadratic Inequalities.* To solve the constraints generated in the previous step, we directly integrate the technique of [1] into our algorithm. This is a sound and semi-complete approach based on Putinar's Positivstellensatz. We will provide an example below, but refer to [1] for technical details and proofs of soundness/completeness of this step.

In this step, for each constraint of the form  $\Phi \Rightarrow \Psi$ , the algorithm first rewrites

<sup>1</sup> We have to find values for  $c$ -variables that satisfy all these constraints conjunctively. This is why we have an extra existential quantifier in (1).

$\Phi$  in disjunctive normal form as  $\phi_1 \vee \dots \vee \phi_t$  and  $\Psi$  in conjunctive normal form as  $\Psi \equiv \psi_1 \wedge \dots \wedge \psi_r$ . Then for each  $1 \leq i \leq t$  and  $1 \leq j \leq r$  the algorithm uses Putinar's Positivstellensatz in the exact same way as in [1] to generate a set of quadratic inequalities equivalent to  $\phi_i \Rightarrow \psi_j$ . The algorithm keeps track of a quadratic program  $\Gamma$  and adds these new inequalities to it conjunctively.

4. *Handling Initial Conditions.* Additionally, for every variable  $x \in \mathcal{V}$ , the algorithm introduces another symbolic template variable  $t_x$ , modeling the initial value of  $x$  in the program, and adds the constraint  $[\theta_{init}(t) \wedge f_{l_{init}}(t) \geq 0]$  to  $\Gamma$  to impose that there exists an initial state in  $\mathcal{T}$  at which the value of the EBRF  $f$  is non-negative.
5. *Solving the System.* Finally, the algorithm uses an external solver (usually an SMT solver) to compute values of  $t$  and  $c$  variables for which  $\Gamma$  is satisfied. If the solver succeeds in solving the system of constraints  $\Gamma$ , the computed values of  $c$  and  $t$  variables give rise to a concrete instance of an  $\mathcal{B}$ -EBRF for  $\mathcal{T}$ . This implies that the answer to the EB-PA problem is positive, and the algorithm return “Yes”. Otherwise, the algorithm returns “Unknown”, as there might exist a  $\mathcal{B}$ -EBRF for  $\mathcal{T}$  of higher maximum polynomial degree  $D$  or a non-polynomial  $\mathcal{B}$ -EBRF.

**Theorem 3 (Existential Soundness and Semi-completeness).** *The algorithm above is a sound and semi-complete reduction to quadratic programming for synthesizing an EBRF in a polynomial transition system  $\mathcal{T}$  given a Büchi specification  $\mathcal{B}$  obtained from an LTL formula with polynomial inequalities in AP. Moreover, for any fixed  $D$ , the algorithm has sub-exponential complexity.*

In the above theorem, soundness means that every solution to the QP instance is a valid EBRF and semi-completeness means that if a polynomial EBRF exists and the chosen maximum degree  $D$  is large enough, then the QP instance will have a solution. In practice, we simply pass the QP instance to an SMT solver. Since it does not include a quantifier alternation, the SMT solvers have dedicated heuristics and are quite efficient on QP instances.

**Synthesis of Polynomial UBRFs.** Our algorithm for synthesis of UBRFs is almost the same as our EBRF algorithm, except that the constraints generated in Steps 2 and 4 are slightly different.

**Changes to Step 2.** Step 2 is the main difference between the two algorithms. In this step, for each location  $l \in L$  and each transition  $\tau \in Out_l$  the UBRF algorithm adds  $(\phi_{l,\tau} \Rightarrow \psi_{l,\tau})$  to the set of constraints, where we have  $\phi_{l,\tau} \equiv \theta_l \wedge G_\tau \wedge f_l(x) \geq 0$  and  $\psi_{l,\tau} \equiv \mathcal{B}\text{-Rank}(\tau)$ . The intuition behind this step is that whenever a transition is enabled, it has to be Büchi-ranked by  $f$ .

**Changes to Step 4.** In this step, instead of searching for a suitable initial valuation for program variables, the algorithm adds the quadratic inequalities equivalent to  $(\theta_{init} \Rightarrow f_{l_{init}}(x) \geq 0)$  to  $\Gamma$ . The quadratic inequalities are obtained exactly as in Step 3. This is because the value of the UBRF must be non-negative on every initial state of the transition system.

In the universal case, we have a similar theorem of soundness and semi-completeness whose proof is exactly the same as Theorem 3.

**Theorem 4 (Universal Soundness and Semi-completeness).** *The algorithm above is a sound and semi-complete reduction to quadratic programming for synthesizing an UBRF in a polynomial transition system  $\mathcal{T}$  given a Büchi specification  $\mathcal{B}$  obtained from an LTL formula with polynomial inequalities in AP. Moreover, for any fixed maximum polynomial degree  $D$ , the algorithm has sub-exponential complexity.*

## 5 Experimental Results

**General Setup of Experiments.** We implemented a prototype<sup>2</sup> of our UBRF and EBRF synthesis algorithms in Java and used Z3 [50], Barcelogic [6] and MathSAT5 [21] to solve the generated systems of quadratic inequalities. More specifically, after obtaining the QP instance, our tool calls all three SMT solvers in parallel. We also used ASPIC [34] for invariant generation for benchmarks that are linear programs. Experiments were performed on a Debian 11 machine with a 2.60GHz Intel E5-2670 CPU and 6 GB of RAM with a timeout of 1800s.

**Baselines.** We compare our tool with Ultimate LTLAutomizer [31], nuXmv [20], and MuVal [65] as well as with a modification of our method that instead of using Putinar’s Positivstellensatz simply passes entailment constraints to the SMT-solver Z3 [50]:

- Ultimate LTLAutomizer makes use of “Büchi programs”, which is a similar notion to our product of a transition system and a Büchi Automaton, to either prove that every lasso shaped path in the input program satisfies the given LTL formula, or find a path that violates it. However, in contrast to our tool, it neither supports non-linear programs nor provides completeness.
- nuXmv is a symbolic model checker with support for finite and infinite transition systems. It allows both existential and universal LTL program analysis and supports non-linear programs. It does not provide any completeness guarantees.
- MuVal [65] is a fixed-point logic validity checker based on pfwCSP solving [66]. It supports both linear and non-linear programs with integer variables and recursive functions.
- When directly applying Z3, instead of the dedicated quantifier elimination method (Step 3 of our algorithm), we directly pass the quantified formula (1) to the solver, which will in turn apply its own generic quantifier elimination. This is an ablation experiment to check whether Step 3 is needed in practice.

**Benchmarks.** We gathered benchmarks from two sources:

- 297 benchmarks from the “Termination of C-Integer Programs” category of TermComp’22 [37]<sup>3</sup>. Among these, 287 programs only contained linear arith-

<sup>2</sup> Available at [github.com/ekgma/LTL-VerP](https://github.com/ekgma/LTL-VerP).

<sup>3</sup> There were originally 335 benchmarks, but we had to remove benchmarks with unbounded non-determinism and those without any variables, since they cannot be translated to transition systems and are not supported in our setting.

metic which is supported by all comparator tools, whereas 10 programs contained polynomial expressions not supported by Ultimate.

- 21 non-linear benchmarks from the “ReachSafety-Loops `nl-digbench`” category of SV-COMP’22 [5]<sup>4</sup>. As these benchmarks are all non-linear, none of them are supported by Ultimate.

**LTL Specifications.** We used the four LTL specifications shown in Table 1. In all four considered specifications,  $x$  represents the alphabetically first variable in the input program. The motivation behind our specifications is as follows:

- *Reach-Avoid (RA) Specifications.* The first specification is an example of a reach-avoid specification, which specifies that a program run should terminate without ever making  $x$  negative. Reach-avoid specifications are standard in the analysis of dynamical and hybrid systems [48, 61, 67]. Another example is requiring a program to termination while satisfying all program assertions.
- *Overflow (OV) Specifications.* Intuitively, we want to evaluate whether our approach is capable of detecting variable overflows. The second specification specifies that each program run either terminates or the value of the variable  $x$  overflows. Specifically, suppose that an overflow is handled as a runtime error and ends the program. The negation (refutation) of this specification models the existence of a run that neither terminates nor overflows and so converges.
- *Recurrence (RC) Specifications.* The third specification is an instance of recurrence specifications which specify that a program run visits a set of states infinitely many times [47]. Our example requires that a program run contains infinitely many visits to states in which  $x$  has a non-negative value.
- *Progress (PR) Specifications.* The fourth specification is an example of progress specifications. In our experimental evaluation, progress specification specifies that a program run always makes progress from states in which the value of  $x$  is less than  $-5$  to states in which the value of  $x$  is strictly positive.

**Table 1.** LTL specifications used in our experiments.

Name	Formula	Pre-condition $\theta_{init}$
<i>RA</i>	$(F \text{ at}(l_{term})) \wedge (G x \geq 0)$	$\forall x \in \mathcal{V}, 0 \leq x \leq 64$
<i>OV</i>	$F (\text{at}(l_{term}) \vee x < -64 \vee x > 63)$	$\forall x \in \mathcal{V}, -64 \leq x \leq 63$
<i>RC</i>	$G F (x \geq 0)$	$\forall x \in \mathcal{V}, -64 \leq x \leq 63$
<i>PR</i>	$G (x < -5 \Rightarrow F (x > 0))$	$\forall x \in \mathcal{V}, -64 \leq x \leq 63$

**Results on Linear Programs.** The top rows of Table 2 summarize our results over linear benchmarks to which all tools are applicable. First, we observe that in all cases our tool outperforms the method that uses Z3 for quantifier elimination, showing that our Step 3 is a crucial and helpful part of the algorithm. Compared

<sup>4</sup> The original benchmark set contains 28 programs, but 7 of them contain unsupported operators such as integer mod and are thus not expressible in our setting.

**Table 2.** Summary of our experimental results. For each class of benchmarks (linear/non-linear) and each formula, We report in how many cases the tool could successfully prove the formula (Yes) or refute it (No), total number of cases proved by the tool (Tot.), number of instances uniquely solved by each tool and no other tools (U.), and average runtime of each tool on programs that were successfully proved as correct with respect to each specification (Avg. T).

Formula		Ours				Ultimate				nuXmv				MuVal				Z3			
		Yes	No	Tot.	U.	Yes	No	Tot.	U.	Yes	No	Tot.	U.	Yes	No	Tot.	U.	Yes	No	Tot.	U.
Linear	RA	141	114	255	5	142	121	263	7	76	91	137	0	118	76	194	0	56	36	92	0
	OV	199	47	246	4	212	55	267	5	110	50	160	0	205	47	252	3	48	27	75	0
	RC	87	187	274	0	86	194	280	0	83	183	266	0	86	191	277	0	44	71	115	0
	PR	43	222	265	1	45	237	282	0	44	227	271	0	42	235	277	0	29	77	106	0
	Avg. T	5.4	81.5	47.2	-	5.4	4.1	4.7	-	248.9	13.5	98.7	-	48.8	8.43	26.4	-	18.5	160.6	95.7	-
Non-linear	RA	24	3	27	8	-	-	-	-	1	0	1	0	18	1	19	2	0	0	0	0
	OV	26	0	26	2	-	-	-	-	7	0	7	0	25	0	25	1	0	0	0	0
	RC	20	6	26	0	-	-	-	-	17	9	26	2	17	7	24	2	0	0	0	0
	PR	11	16	27	1	-	-	-	-	9	16	25	0	5	16	21	1	0	0	0	0
	Avg. T	10.7	99.1	32.3	-	-	-	-	-	34.6	0.3	20.0	-	109.6	14.7	84.7	-	-	-	-	-

to nuXmv, our tool proves more instances in all but two LTL refutation and one LTL verification cases, i.e. the “No” column for the OV and PR specifications and the “Yes” column for the PR specification. On the other hand, our prototype tool is on par with Ultimate and MuVal, while proving 10 unique instances. Note that Ultimate is a state of the art and well-maintained competition tool that is highly optimized with heuristics that aim at the linear case. In contrast, it cannot handle polynomial instances. Our results shown in Table 2 demonstrate that our prototype tool is very competitive already on linear benchmarks, even though our main contribution is to provide practically-efficient semi-complete algorithms for the polynomial case.

**Unique Instances.** An important observation is that our tool successfully handles 10 unique *linear* instances that no other tool manages to prove or refute. Thus, our evaluation shows that our method handles not only polynomial, but even linear benchmarks that were beyond the reach of the existing methods. This shows that our algorithm, besides the desired theoretical guarantee of semi-completeness, provides an effective automated method. Future advances in invariant generation and SMT solving will likely further improve the performance.

**Runtimes.** Our tool and Ultimate are the fastest tools for proving LTL verification instances with an equal average runtime of 5.4 s. For LTL refutation, our tool is slower than other tools.

**Results on Non-linear Programs.** The bottom rows of Table 2 show the performance of our tool and the baselines on the non-linear benchmarks. Ultimate does not support non-linear arithmetic and Z3 timed out on every benchmark in this category. Here, compared to nuXmv, our tool succeeded in solving strictly more instances in all but one formula, i.e. *RC*, where both tools solve the same number of instances. In comparison with MuVal, our tool proves more instances for all four formulas. Moreover, the fact that Z3 timed out for every program in this table is further confirmation of the practical necessity of Step 3 (Quantifier

Elimination Procedure of [1]) in our algorithm. Note that our prototype could prove 11 instances that none of the other tools could handle.

**Summary.** Our experiments demonstrate that our automated algorithms are able to synthesize both LTL verification and refutation witnesses for a wide variety of programs. Our technique outperforms the previous methods when given non-linear polynomial programs (Bottom rows of Table 2). Moreover, even in the much more widely-studied case of linear programs, we are able to handle instances that were beyond the reach of previous methods and to solve the number of instances that is close to the state-of-the-art tools (Top Rows of Table 2).

## 6 Conclusion

We presented a novel family of sound and complete witnesses for template-based LTL verification. Our approach is applicable to both verification and refutation of LTL properties in programs. It unifies and significantly generalizes previous works targeting special cases of LTL, e.g. termination, safety and reachability. We also showed that our LTL witnesses can be synthesized in a sound and semi-complete manner by a reduction to quadratic programming. Our reduction works when the program and the witness are both polynomial. An interesting direction of future work would be to consider non-numerical programs that allow heap-manipulating operations. A common approach to handling heap-manipulating operations is to construct numerical abstractions of programs [7, 46] and perform the analysis on numerical abstractions. Thus, coupling such approaches, e.g. [28], with our method is a compelling future direction.

**Acknowledgements.** This work was supported in part by the ERC-2020-CoG 863818 (FoRM-SMArt) and the Hong Kong Research Grants Council ECS Project Number 26208122.

**Data Availability Statement.** The implementations of the algorithms mentioned in the experiments section and the benchmarks are available at [doi.org/10.5281/zenodo.12518217](https://doi.org/10.5281/zenodo.12518217).

## References

1. Asadi, A., Chatterjee, K., Fu, H., Goharshady, A.K., Mahdavi, M.: Polynomial reachability witnesses via stellensätze. In: PLDI, pp. 772–787 (2021)
2. Baier, C., Katoen, J.: Principles of Model Checking. MIT Press (2008)
3. Baresi, L., Kallehbasti, M.M.P., Rossi, M.: Efficient scalable verification of LTL specifications. In: ICSE (1), pp. 711–721. IEEE Computer Society (2015)
4. Bauch, P., Havel, V., Barnat, J.: LTL model checking of LLVM bitcode with symbolic data. In: MEMICS, pp. 47–59 (2014)
5. Beyer, D.: Progress on software verification: SV-COMP 2022. In: TACAS, pp. 375–402 (2022)



6. Bofill, M., Nieuwenhuis, R., Oliveras, A., Rodríguez-Carbonell, E., Rubio, A.: The barcelogic SMT solver. In: Gupta, A., Malik, S. (eds.) CAV 2008. LNCS, vol. 5123, pp. 294–298. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-70545-1\\_27](https://doi.org/10.1007/978-3-540-70545-1_27)
7. Bouajjani, A., Bozga, M., Habermehl, P., Iosif, R., Moro, P., Vojnar, T.: Programs with lists are counter automata. *Formal Methods Syst. Des.* **38**(2), 158–192 (2011)
8. Brockschmidt, M., Cook, B., Ishtiaq, S., Khlaaf, H., Piterman, N.: T2: temporal property verification. In: Chechik, M., Raskin, J.-F. (eds.) TACAS 2016. LNCS, vol. 9636, pp. 387–393. Springer, Heidelberg (2016). [https://doi.org/10.1007/978-3-662-49674-9\\_22](https://doi.org/10.1007/978-3-662-49674-9_22)
9. Büchi, J.R.: Symposium on decision problems: on a decision method in restricted second order arithmetic. In: *Studies in Logic and the Foundations of Mathematics*, vol. 44, pp. 1–11 (1966)
10. Cai, Z., Farokhnia, S., Goharshady, A.K., Hitarth, S.: Asparagus: automated synthesis of parametric gas upper-bounds for smart contracts. In: OOPSLA (2023)
11. Chatterjee, K., Fu, H., Goharshady, A.K.: Termination analysis of probabilistic programs through positivstellensatz’s. In: CAV, pp. 3–22 (2016)
12. Chatterjee, K., Fu, H., Goharshady, A.K., Goharshady, E.K.: Polynomial invariant generation for non-deterministic recursive programs. In: PLDI, pp. 672–687 (2020)
13. Chatterjee, K., Goharshady, A.K., Goharshady, E.K., Karrabi, M., Zikelic, D.: Sound and complete witnesses for template-based verification of LTL properties on polynomial programs. *arXiv preprint arXiv:2403.05386* (2024)
14. Chatterjee, K., Goharshady, A.K., Meggendorfer, T., Zikelic, D.: Quantitative bounds on resource usage of probabilistic programs. In: OOPSLA (2024)
15. Chatterjee, K., Goharshady, A.K., Meggendorfer, T., Zikelic, D.: Sound and complete certificates for quantitative termination analysis of probabilistic programs. In: CAV, pp. 55–78 (2022)
16. Chatterjee, K., Goharshady, E.K., Novotný, P., Žikelić, D.: Proving non-termination by program reversal. In: PLDI, pp. 1033–1048 (2021)
17. Chatterjee, K., Goharshady, E.K., Novotný, P., Žikelić, U.: Equivalence and similarity refutation for probabilistic programs (PLDI) (2024). <https://doi.org/10.1145/3656462>
18. Chatterjee, K., Novotný, P., Žikelić, D.: Stochastic invariants for probabilistic termination. In: Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, 18–20 January 2017, pp. 145–160. ACM (2017). <https://doi.org/10.1145/3009837.3009873>
19. Cimatti, A., Griggio, A., Magnago, E.: LTL falsification in infinite-state systems. *Inf. Comput.* **289**, 104977 (2022)
20. Cimatti, A., Griggio, A., Magnago, E., Roveri, M., Tonetta, S.: Extending nuXmv with timed transition systems and timed temporal properties. In: Dillig, I., Tasiran, S. (eds.) CAV 2019. LNCS, vol. 11561, pp. 376–386. Springer, Cham (2019). [https://doi.org/10.1007/978-3-030-25540-4\\_21](https://doi.org/10.1007/978-3-030-25540-4_21)
21. Cimatti, A., Griggio, A., Schaafsma, B.J., Sebastiani, R.: The mathsat5 SMT solver. In: TACAS, pp. 93–107 (2013)
22. Clark, A.: Verification and synthesis of control barrier functions. In: CDC, pp. 6105–6112 (2021)
23. Clarke, E.M., Grumberg, O., Jha, S., Lu, Y., Veith, H.: Counterexample-guided abstraction refinement. In: CAV (2000)
24. Clarke, E.M., Henzinger, T.A., Veith, H., Bloem, R.: *Handbook of Model Checking*. Springer (2018)

25. Colón, M.A., Sankaranarayanan, S., Sipma, H.B.: Linear invariant generation using non-linear constraint solving. In: CAV, pp. 420–432 (2003)
26. Cook, B., Khlaaf, H., Piterman, N.: Fairness for infinite-state systems. In: Baier, C., Tinelli, C. (eds.) TACAS 2015. LNCS, vol. 9035, pp. 384–398. Springer, Heidelberg (2015). [https://doi.org/10.1007/978-3-662-46681-0\\_30](https://doi.org/10.1007/978-3-662-46681-0_30)
27. Cook, B., Koskinen, E.: Making prophecies with decision predicates. In: POPL, pp. 399–410 (2011)
28. Cook, B., Koskinen, E.: Reasoning about nondeterminism in programs. In: PLDI, pp. 219–230 (2013)
29. Daniel, J., Cimatti, A., Griggio, A., Tonetta, S., Mover, S.: Infinite-state liveness-to-safety via implicit abstraction and well-founded relations. In: Chaudhuri, S., Farzan, A. (eds.) CAV 2016. LNCS, vol. 9779, pp. 271–291. Springer, Cham (2016). [https://doi.org/10.1007/978-3-319-41528-4\\_15](https://doi.org/10.1007/978-3-319-41528-4_15)
30. De Branges, L.: The Stone-Weierstrass theorem. *Proc. AMS* **10**(5), 822–824 (1959)
31. Dietsch, D., Heizmann, M., Langenfeld, V., Podelski, A.: Fairness modulo theory: a new approach to LTL software model checking. In: CAV, pp. 49–66 (2015)
32. Farkas, J.: Theorie der einfachen ungleichungen. *Journal für die reine und angewandte Mathematik* **1902**(124), 1–27 (1902)
33. Farzan, A., Kincaid, Z., Podelski, A.: Proving liveness of parameterized programs. In: LICS, pp. 185–196 (2016)
34. Feautrier, P., Gonnord, L.: Accelerated invariant generation for C programs with *aspic* and *c2fsm*. *Electron. Notes Theor. Comput. Sci.* 3–13 (2010)
35. Feng, Y., Zhang, L., Jansen, D.N., Zhan, N., Xia, B.: Finding polynomial loop invariants for probabilistic programs. In: ATVA, pp. 400–416 (2017)
36. Floyd, R.W.: Assigning meanings to programs. In: *Program Verification: Fundamental Issues in Computer Science*, pp. 65–81 (1993)
37. Frohn, F., Giesl, J., Moser, G., Rubio, A., Yamada, A., et al.: Termination competition 2022 (2021). [https://termination-portal.org/wiki/Termination\\_Competition\\_2022](https://termination-portal.org/wiki/Termination_Competition_2022)
38. Fulton, N.: Verifiably safe autonomy for cyber-physical systems. Ph.D. thesis, Carnegie Mellon University (2018)
39. Funke, F., Jantsch, S., Baier, C.: Farkas certificates and minimal witnesses for probabilistic reachability constraints. In: TACAS, pp. 324–345 (2020)
40. Graf, S., Saïdi, H.: Construction of abstract state graphs with PVS. In: CAV, pp. 72–83 (1997)
41. Gulwani, S., Srivastava, S., Venkatesan, R.: Program analysis as constraint solving. In: PLDI, pp. 281–292 (2008)
42. Gurriet, T., Singletary, A., Reher, J., Ciarletta, L., Feron, E., Ames, A.D.: Towards a framework for realizable safety critical control through active set invariance. In: ICCPS, pp. 98–106 (2018)
43. Heizmann, M., Hoenicke, J., Leike, J., Podelski, A.: Linear ranking for linear Lasso programs. In: Van Hung, D., Ogawa, M. (eds.) ATVA 2013. LNCS, vol. 8172, pp. 365–380. Springer, Cham (2013). [https://doi.org/10.1007/978-3-319-02444-8\\_26](https://doi.org/10.1007/978-3-319-02444-8_26)
44. Huang, M., Fu, H., Chatterjee, K., Goharshady, A.K.: Modular verification for almost-sure termination of probabilistic programs. *Proc. ACM Program. Lang.* **3**(OOPSLA), 129:1–129:29 (2019)
45. Kincaid, Z., Cyphert, J., Breck, J., Reps, T.W.: Non-linear reasoning for invariant synthesis. In: POPL, pp. 54:1–54:33 (2018)
46. Magill, S., Tsai, M., Lee, P., Tsay, Y.: Automatic numeric abstractions for heap-manipulating programs. In: POPL, pp. 211–222 (2010)

47. Manna, Z., Pnueli, A.: A hierarchy of temporal properties. In: PODC, pp. 377–410 (1990)
48. Meng, Y., Liu, J.: Lyapunov-barrier characterization of robust reach-avoid-stay specifications for hybrid systems (2022). <https://doi.org/10.48550/ARXIV.2211.00814>
49. Moosbrugger, M., Bartocci, E., Katoen, J.-P., Kovács, L.: The probabilistic termination tool Amber. In: Huisman, M., Păsăreanu, C., Zhan, N. (eds.) FM 2021. LNCS, vol. 13047, pp. 667–675. Springer, Cham (2021). [https://doi.org/10.1007/978-3-030-90870-6\\_36](https://doi.org/10.1007/978-3-030-90870-6_36)
50. de Moura, L., Bjørner, N.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008). [https://doi.org/10.1007/978-3-540-78800-3\\_24](https://doi.org/10.1007/978-3-540-78800-3_24)
51. Neumann, E., Ouaknine, J., Worrell, J.: On ranking function synthesis and termination for polynomial programs. In: CONCUR, pp. 15:1–15:15 (2020)
52. Padon, O., Hoenicke, J., McMillan, K.L., Podelski, A., Sagiv, M., Shoham, S.: Temporal prophecy for proving temporal properties of infinite-state systems. *Formal Methods Syst. Des.* **57**(2), 246–269 (2021)
53. Pnueli, A.: The temporal logic of programs. In: FOCS, pp. 46–57 (1977)
54. Podelski, A., Rybalchenko, A.: A complete method for the synthesis of linear ranking functions. In: Steffen, B., Levi, G. (eds.) VMCAI 2004. LNCS, vol. 2937, pp. 239–251. Springer, Heidelberg (2004). [https://doi.org/10.1007/978-3-540-24622-0\\_20](https://doi.org/10.1007/978-3-540-24622-0_20)
55. Podelski, A., Rybalchenko, A.: Transition predicate abstraction and fair termination. In: POPL, pp. 132–144 (2005)
56. Rice, H.G.: Classes of recursively enumerable sets and their decision problems. *Trans. AMS* **74**(2), 358–366 (1953)
57. Sankaranarayanan, S., Sipma, H., Manna, Z.: Non-linear loop invariant generation using gröbner bases. In: POPL, pp. 318–329 (2004)
58. Sankaranarayanan, S., Sipma, H.B., Manna, Z.: Constraint-based linear-relations analysis. In: Giacobazzi, R. (ed.) SAS 2004. LNCS, vol. 3148, pp. 53–68. Springer, Heidelberg (2004). [https://doi.org/10.1007/978-3-540-27864-1\\_7](https://doi.org/10.1007/978-3-540-27864-1_7)
59. Shen, L., Wu, M., Yang, Z., Zeng, Z.: Generating exact nonlinear ranking functions by symbolic-numeric hybrid method. *J. Syst. Sci. Complex.* **26**(2), 291–301 (2013)
60. Strejcek, J.: Linear temporal logic: expressiveness and model checking. Ph.D. thesis, Masaryk University (2004)
61. Summers, S., Lygeros, J.: Verification of discrete time stochastic hybrid systems: a stochastic reach-avoid decision problem. *Autom.* 1951–1961 (2010)
62. Sun, Y., Fu, H., Chatterjee, K., Goharshady, A.K.: Automated tail bound analysis for probabilistic recurrence relations. In: CAV, pp. 16–39 (2023)
63. Takisaka, T., Oyabu, Y., Urabe, N., Hasuo, I.: Ranking and repulsing supermartingales for reachability in randomized programs. *TOPLAS* **43**(2), 5:1–5:46 (2021)
64. Turing, A.M.: On computable numbers, with an application to the entscheidungsproblem. *J. Math.* **58**(345–363), 5 (1936)
65. Unno, H., Terauchi, T., Gu, Y., Koskinen, E.: Modular primal-dual fixpoint logic solving for temporal verification. In: POPL, pp. 2111–2140 (2023)
66. Unno, H., Terauchi, T., Koskinen, E.: Constraint-based relational verification. In: CAV, pp. 742–766 (2021)
67. Žikelić, D., Lechner, M., Henzinger, T.A., Chatterjee, K.: Learning control policies for stochastic systems with reach-avoid guarantees. In: AAI, pp. 11926–11935 (2023)

68. Wang, J., Sun, Y., Fu, H., Chatterjee, K., Goharshady, A.K.: Quantitative analysis of assertion violations in probabilistic programs. In: PLDI, pp. 1171–1186 (2021)
69. Wang, P., Fu, H., Goharshady, A.K., Chatterjee, K., Qin, X., Shi, W.: Cost analysis of nondeterministic probabilistic programs. In: PLDI, pp. 204–220 (2019)
70. Wang, Q., Chen, M., Xue, B., Zhan, N., Katoen, J.: Synthesizing invariant barrier certificates via difference-of-convex programming. In: CAV, pp. 443–466 (2021)
71. Zhang, Y., Yang, Z., Lin, W., Zhu, H., Chen, X., Li, X.: Safety verification of nonlinear hybrid systems based on bilinear programming. *IEEE Trans. Comput. Aided Des. Integr. Circuits Syst.* **37**(11), 2768–2778 (2018)

**Open Access** This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter's Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter's Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.

