

Temporal Prophecy for Proving Temporal Properties of Infinite-State Systems

Oded Padon^{*}, Jochen Hoenicke[†], Kenneth L. McMillan[‡], Andreas Podelski[†], Mooly Sagiv^{*} and Sharon Shoham^{*}
^{*}Tel Aviv University, Israel [†]University of Freiburg, Germany [‡]Microsoft Research, USA

Abstract—Various verification techniques for temporal properties transform temporal verification to safety verification. For infinite-state systems, these transformations are inherently imprecise. That is, for some instances, the temporal property holds, but the resulting safety property does not. This paper introduces a mechanism for tackling this imprecision. This mechanism, which we call *temporal prophecy*, is inspired by prophecy variables. Temporal prophecy refines an infinite-state system using first-order linear temporal logic formulas, via a suitable tableau construction. For a specific liveness-to-safety transformation based on first-order logic, we show that using temporal prophecy strictly increases the precision. Furthermore, temporal prophecy leads to robustness of the proof method, which is manifested by a cut elimination theorem. We integrate our approach into the Ivy deductive verification system, and show that it can handle challenging temporal verification examples.

1. Introduction

There are various techniques in the literature that transform the problem of verifying liveness of a system to the problem of verifying safety of a different system. These transformations compose the system with a device that has the known property that some safety condition σ implies liveness. The classical example of this is proving termination of a while loop with a ranking function. In this case, the device evaluates a chosen function r on loop entry, where the range of r is a well-founded set. The safety property σ is that r decreases at every iteration, which implies that the loop must terminate.

A related transformation, due to Armin Biere [5], applies to finite-state (possibly parameterized) systems. The safety property σ is, in effect, that no state occurs twice, from which we can infer termination. In the infinite-state case, this can be generalized using a function f that projects the program state onto a finite set. We can think of this as a ranking that tracks the set of unseen values of f and is ordered by set inclusion. However, the property that no value of f occurs twice is simpler to verify, since the composed device can non-deterministically guess the recurring value. In general, the effectiveness of a liveness-to-safety transformation depends strongly on the difficulty of the resulting safety proof problem.

Other methods can be seen as instances of this general approach. For example, the Terminator tool [11] might be seen as combining the ranking and the finite projection

approaches. Another approach by Fang *et al.* applies a collection of ad-hoc devices with known safety-to-liveness properties to prove liveness of parameterized protocols [16]. Of greatest interest here, a recent paper by Padon *et al.* uses a dynamically chosen finite projection that depends on a finite prefix of the system’s execution [30]. The approach of [28] also has some similar characteristics.

In the case of infinite-state systems, these transformations from liveness verification to safety verification are not precise reductions. That is, while safety implies liveness, a counterexample to the safety property σ does not in general imply a counterexample to liveness. For example, in the projection method, a terminating infinite-state system may have runs whose length exceeds the finite range of any chosen projection f , forcing some value to repeat.

In this paper, we show that the precision of a liveness-to-safety transformation can be usefully increased by the addition of *prophecy variables*. These variables are expressed as first-order LTL formulas. For example, suppose we augment the state of the system with a variable $r_{\Box p}$ that tracks the truth value of the proposition $\Box p$, which is true when p holds in all future states. We can soundly add two constraints to the transition system. To the transition relation, we add $r_{\Box p} \leftrightarrow (p \wedge r_{\Box p}')$, where $r_{\Box p}'$ denotes the value of the prophecy variable in the post-state. We also add the fairness constraint that $r_{\Box p} \vee \neg p$ holds infinitely often. These constraints are typical of tableau constructions that convert a temporal formula to a symbolic automaton. As we show in this paper, the additional information they provide refines the trace set of the transformed system, potentially eliminating false counterexamples.

In particular, we will show how to integrate tableau-based prophecy with the liveness-to-safety transformation of [30] that uses a history-based finite projection, referred to as *dynamic abstraction*. We show that the precision of this transformation is consequently increased. The result is that we can prove properties that otherwise would not be directly provable using the technique.

This paper makes the following contributions:

- 1) Introduce the notion of temporal prophecy, including prophecy formulas and prophecy witnesses, *via* a first-order LTL tableau construction.
- 2) Show that temporal prophecy increases the proof power (i.e., precision) of the safety-to-liveness transformation based on dynamic abstraction, and further show that the properties provable with tem-




Figure 1. The ticket mutual exclusion protocol. Edges are labeled by condition / action.

poral prophecy are closed under *first-order reasoning*, with cut elimination as a special case.

- 3) Integrate the liveness-to-safety transformation based on dynamic abstraction and temporal prophecy into the Ivy deductive verification system, deriving the prophecy formulas from an inductive invariant provided by the user (for proving the safety property).
- 4) Demonstrate the effectiveness of the approach on some challenging examples that cannot be handled by the transformation without temporal prophecy.
- 5) Demonstrate that prophecy witnesses can eliminate quantifier alternations in the verification conditions generated for the safety problem obtained after the transformation, facilitating decidable reasoning.

2. Illustrative Example

We illustrate our approach using the ticket protocol for ensuring mutual exclusion with non starvation among multiple threads, depicted in Fig. 1. The ticket protocol may be run by any number of threads, and also allows dynamic spawning of threads. The protocol is an idealized version of spinlocks used in the Linux kernel [13]. In the protocol, each thread can be in one of three states: idle, waiting to enter the critical section, or in the critical section. The right to enter the critical section is determined by a ticket number. A global variable n , records the next available ticket, and a global variable s , records the ticket currently being served. Each thread has a local variable m that records the ticket it holds. A thread only enters the critical section when $m \leq s$. Once a thread enters the critical section, it handles tasks that accumulated in its task queue, and stays in the critical section until its queue is empty (tasks are only added to the queue when the thread is outside the critical section). In Fig. 1, this is modeled by the task counter q , a thread-local variable which is non-deterministically set when a thread enters the critical section (to account for the unbounded, but finite, number of tasks), and is then decremented in each step. When $q = 0$ the thread leaves the critical section, and increments s to allow other threads to be served.

The protocol is designed to satisfy the following first-order temporal property:

$$(\forall x. \Box \Diamond \text{scheduled}(x)) \rightarrow \forall y. \Box (\text{wait}(y) \rightarrow \Diamond \text{critical}(y))$$

That is, if every process is scheduled infinitely often, then every waiting process eventually enters its critical section. (Note that we encode fairness assumptions as part of the temporal property.)

Insufficiency of liveness-to-safety transformations. While the temporal property is clearly satisfied by the ticket protocol, proving it is challenging for liveness-to-safety transformations. First, due to the unbounded values obtained by the ticket number and the task counter, and also due to dynamic spawning of threads, this example does not belong to the class of parameterized systems [32], where a simple lasso argument is sound (and complete) for proving liveness. Second, while using a finite abstraction can recover soundness, no fixed finite abstraction is precise enough to show the absence of a lasso-shaped counterexample in this example. The reason is that a thread can go to the waiting state (*wait*) with any number of threads waiting “ahead of it in line”.

For cases where no finite abstraction is sufficiently precise to prove liveness, we may instead apply the liveness-to-safety transformation of [30]. This transformation relaxes the requirement of proving absence of lassos over a fixed finite abstraction, and instead requires one to prove absence of lassos over a *dynamic* finite abstraction that is only determined after some prefix of the trace (allowing for better precision). Soundness is maintained since the abstraction is still finite. Technically, the technique requires to prove that no *abstract lasso* exists, where an abstract lasso is a finite execution prefix that (i) visits a *freeze point*, at which a finite projection (abstraction) of the state space is fixed, (ii) the freeze point is followed by two states that are equal in the projection. We refer to these as the *repeating states*, and (iii) all fairness constraints are visited both before the freeze point and between the repeating states.

Unlike fixed finite abstractions, dynamic abstractions allow us to prove that an eventuality holds if there is a finite upper bound on the number of steps required *at the time the eventuality is asserted* (the freeze point). The bound need not be fixed *a priori*. Unfortunately, due to the non-determinism introduced by the task counter q , each of the k threads ahead of t in line could require an unbounded number of steps to leave the critical section, and this number is not yet determined when t makes its request. As a result, there is an abstract lasso which freezes the abstraction when t makes its request, after which some other thread t_0 enters the critical section and loops, decrementing its task counter q . Since the value of the task counter of t_0 is not captured in the abstraction, the loop does not change the abstract state. This spurious abstract lasso prevents this liveness-to-safety transformation from proving the property.

Temporal prophecy to the rescue. The key to fixing this problem is to predict the future to the extent that a bound on the steps required for progress is determined at the freeze point. Surprisingly, this is accomplished by the use of one temporal prophecy variable corresponding to the truth value of the following formula:

$$\exists x. \Diamond \Box \text{critical}(x).$$

If this formula is initially true, there is some thread t_0 that eventually enters the critical section and stays there. At this point, we can prove it eventually exits (a contradiction)

because the number of steps needed for this is bounded by the current task counter of t_0 . Operationally, the freeze point is delayed until $\Box \text{critical}(x)$ holds at which point t_0 's task counter is captured in the finite projection, ruling out an abstract lasso. On the other hand if the prophecy variable is initially false, then all threads are infinitely often out of the critical section. With this fairness constraint, thread t requires only a finite number of steps to be served, determined by the number of threads with lesser tickets. Operationally, the extra fairness constraint extends the lasso loop until the abstract state must change, ruling out an abstract lasso.

Though the liveness-to-safety transformation via dynamic abstraction and abstract lasso detection cannot handle the problem as given, introducing suitable temporal prophecy eliminates the spurious abstract lassos. Some spurious lassos are eliminated by postponing the freeze point, thus refining the finite abstraction, and others are eliminated by additional fairness constraints on the lasso loop. This example is explained in greater detail in § 4.3.

3. Preliminaries

In this section, we present the first-order formalism for specifying infinite-state systems and their properties, as well as a tableau construction for first-order LTL formulas.

3.1. Transition Systems in First-Order Logic

A first-order logic transition system is a triple (Σ, ι, τ) , where Σ is a first-order vocabulary that contains only relation symbols and constant symbols (functions can be encoded by relations), ι is a closed formula over Σ defining the set of initial states, and τ is a closed formula over $\Sigma \uplus \Sigma'$, where $\Sigma' = \{\ell' \mid \ell \in \Sigma\}$, defining the transition relation. The constants in Σ represent the program variables.

A state of the transition system is a first-order structure, $s = (\mathcal{D}, \mathcal{I})$, over Σ , where \mathcal{D} denotes the (possibly infinite) domain of the structure and \mathcal{I} denotes the interpretation function. The set of initial states is the set of all states s such that $s \models \iota$, and the set of transitions is the set of all pairs of states (s, s') with the same domain such that $(s, s') \models \tau$. In the latter, (s, s') denotes a structure over the vocabulary $\Sigma \uplus \Sigma'$ with the same domain as s and s' in which the symbols in Σ are interpreted as in s , and the symbols in Σ' are interpreted as in s' .

For a state $s = (\mathcal{D}, \mathcal{I})$ over Σ , and for $D \subseteq \mathcal{D}$, we denote by $s|_D$ the partial structure by projecting s to D , i.e., $s|_D = (D, \mathcal{I}|_D)$, where $\mathcal{I}|_D$ interprets only constants $c \in \Sigma$ for which $\mathcal{I}(c) \in D$ (making it a partial interpretation), and for every relation symbol $r \in \Sigma$ of arity k , $\mathcal{I}|_D(r) = \mathcal{I}(r) \cap D^k$. For a vocabulary $\Sigma' \subseteq \Sigma$, we denote by $s|_{\Sigma'}$ the state over Σ' obtained by restricting the interpretation function to the symbols in Σ' , i.e., $s|_{\Sigma'} = (\mathcal{D}, \mathcal{I}')$, where for every symbol $\ell \in \Sigma'$, $\mathcal{I}'(\ell) = \mathcal{I}(\ell)$.

A (finite or infinite) *trace* of (Σ, ι, τ) is a sequence of states $\pi = s_0, s_1, \dots$ where $s_0 \models \iota$ and $(s_i, s_{i+1}) \models \tau$ for every $0 \leq i < |\pi|$. Every state along the trace has its own

interpretation of the constant and relation symbols, but they all share the same domain.

We note that first-order transition systems are Turing-complete. Furthermore, tools such as Ivy [31] provide modeling languages that are closer to imperative programming languages and compile to a first-order transition system. This makes it easier for a user to provide a first-order specification of the transition system they wish to verify.

Safety. Given a vocabulary Σ , a safety property P is a set of sequences of states over Σ , such that for every sequence of states $\pi \notin P$, there exists a finite prefix π' of π , such that π' and all of its extensions are not in P . A transition system over Σ satisfies P if all of its traces are in P .

3.2. First-Order Linear Temporal Logic (FO-LTL)

To specify temporal properties of first-order transition systems we use First-Order Linear Temporal Logic (FO-LTL), which combines LTL with first-order logic [1]. For simplicity, we consider only the ‘‘globally’’ (\Box) temporal operator. The tableau construction extends to other operators as well, and so does our approach.

Syntax. Given a first-order vocabulary Σ , FO-LTL formulas are defined by:

$$\begin{aligned} f &::= r(t_1, \dots, t_n) \mid t_1 = t_2 \mid \neg f \mid f_1 \vee f_2 \mid \exists x.f \mid \Box f \\ t &::= c \mid x \end{aligned}$$

where r is an n -ary relation symbol in Σ , c is a constant symbol in Σ , x is a variable, each t_i is a term over Σ and \Box denotes the ‘‘globally’’ temporal operator. We also use the standard shorthand for the ‘‘eventually’’ temporal operator: $\Diamond f = \neg \Box \neg f$, and the usual shorthands for logical operators (e.g., $\forall x.f = \neg \exists x.\neg f$).

Semantics. FO-LTL formulas over Σ are interpreted over infinite sequences of states (first-order structures) over Σ . Atomic formulas are interpreted over states, the temporal operators are interpreted as in traditional LTL, and first-order quantifiers are interpreted over the shared domain \mathcal{D} of all states in the trace. Formally, the semantics is defined w.r.t. an infinite sequence of states $\pi = s_0, s_1, \dots$ and an assignment σ that maps variables to \mathcal{D} — the shared domain of all states in π . We define $\pi^i = s_i, s_{i+1}, \dots$ to be the suffix of π starting at index i . The semantics is defined as follows.

$$\begin{aligned} \pi, \sigma &\models r(t_1, \dots, t_n) \Leftrightarrow s_0, \sigma \models r(t_1, \dots, t_n) \\ \pi, \sigma &\models t_1 = t_2 \Leftrightarrow s_0, \sigma \models t_1 = t_2 \\ \pi, \sigma &\models \neg \psi \Leftrightarrow \pi, \sigma \not\models \psi \\ \pi, \sigma &\models \psi_1 \vee \psi_2 \Leftrightarrow \pi, \sigma \models \psi_1 \text{ or } \pi, \sigma \models \psi_2 \\ \pi, \sigma &\models \exists x.\psi \Leftrightarrow \text{exists } d \in \mathcal{D} \text{ s.t. } \pi, \sigma[x \mapsto d] \models \psi \\ \pi, \sigma &\models \Box \psi \Leftrightarrow \text{forall } i \geq 0, \pi^i, \sigma \models \psi \end{aligned}$$

When the formula has no free variables, we omit σ . A first-order transition system (Σ, ι, τ) satisfies a closed FO-LTL formula φ over Σ if all of its traces satisfy φ .

3.3. Tableau for FO-LTL

As part of our liveness-to-safety transformation, we use a standard tableau construction for FO-LTL formulas that results in a first-order transition system with *fairness constraints*. Unlike the classical construction, we define the tableau for a set of formulas, not necessarily a single temporal formula.

For an FO-LTL formula φ , we denote by $sub(\varphi)$ the set of subformulas of φ , defined in the usual way. In the sequel, we consider a finite set A of FO-LTL formulas that is closed under subformulas, i.e. for every $\varphi \in A$, $sub(\varphi) \subseteq A$. Note that A may contain formulas with free variables.

Definition 1 (Tableau vocabulary). Given a finite set A as above over a first-order vocabulary Σ , the *tableau vocabulary* for A , denoted Σ_A , is obtained from Σ by adding a fresh relation symbol $r_{\Box\varphi}$ of arity k for every formula $\Box\varphi \in A$ with k free variables.

Recall that \Box is the only primitive temporal operator we consider (a similar construction can be done for other operators). The symbols added in Σ_A will be used to “label” states by temporal subformulas that are satisfied by all outgoing fair traces. To translate temporal formulas over Σ to first-order formulas over Σ_A we use the following definition.

Definition 2. For a FO-LTL formula $\varphi \in A$ (over Σ), its first-order representation, denoted $FO[\varphi]$, is a first-order formula over Σ_A , defined inductively, as follows.

$$\begin{aligned} FO[\varphi] &= \varphi \quad \text{if } \varphi = r(t_1, \dots, t_n) \text{ or } \varphi = t_1 = t_2 \\ FO[\Box\psi(\bar{x})] &= r_{\Box\psi(\bar{x})}(\bar{x}) \\ FO[\neg\psi] &= \neg FO[\psi] \\ FO[\psi_1 \vee \psi_2] &= FO[\psi_1] \vee FO[\psi_2] \\ FO[\exists x.\psi] &= \exists x.FO[\psi] \end{aligned}$$

Note that $FO[\varphi]$ has the same free variables as φ . We can now define the tableau for A as a transition system.

Definition 3 (Tableau transition system). The *tableau transition system* for A is the first-order transition system $T_A = (\Sigma_A, true, \tau_A)$, where τ_A (defined over $\Sigma_A \uplus \Sigma_A'$) is defined as follows:

$$\tau_A = \bigwedge_{\Box\varphi \in A} \forall \bar{x}. (r_{\Box\varphi}(\bar{x}) \leftrightarrow (FO[\varphi(\bar{x})] \wedge r_{\Box\varphi'}(\bar{x}))).$$

Note that the original symbols in Σ (and Σ') are not constrained by τ_A , and may change arbitrarily with each transition. However, the $r_{\Box\varphi}$ relations are updated in accordance with the property that $\pi, \sigma \models \Box p$ iff $s_0, \sigma \models p$ and $\pi^1, \sigma \models \Box p$ (where s_0 is the first state of π and p is a first-order formula over Σ).

Definition 4 (Fairness). A sequence of states $\pi = s_0, s_1, \dots$ over Σ_A is *A-fair* if for every temporal formula $\Box\varphi(\bar{x}) \in A$ and for every assignment σ , there are infinitely many i 's for which $s_i, \sigma \models FO[\Box\varphi(\bar{x}) \vee \neg\varphi(\bar{x})]$.

Note that $\Box\varphi(\bar{x}) \vee \neg\varphi(\bar{x})$, used above, is equivalent to $\Diamond\neg\varphi(\bar{x}) \rightarrow \neg\varphi(\bar{x})$. So the definition of fairness ensures an eventuality cannot be postponed forever. In the sequel,

the set A is always clear from the context (e.g., from the vocabulary), hence we omit it and simply say that π is fair.

The next claims summarize the properties of the tableau; Lemma 1 states that the FO-LTL formulas over Σ that hold in the outgoing traces of a tableau state correspond to the first-order formulas over Σ_A that hold in the state; Lemma 2 states that every sequence of states over Σ has a representative trace in the tableau; finally, Thm. 1 states that a transition system satisfies a FO-LTL formula iff its product with the tableau of the negated formula has no fair traces.

Lemma 1. In a fair trace $\pi = s_0, s_1, \dots$ of T_A (over Σ_A), for every FO-LTL formula $\psi(\bar{x}) \in A$, for every assignment σ and for every index $i \in \mathbb{N}$, we have that $s_i, \sigma \models FO[\psi(\bar{x})]$ iff $\pi^i, \sigma \models \psi(\bar{x})$.

Lemma 2. Every infinite sequence of states $\hat{s}_0, \hat{s}_1, \dots$ over Σ can be extended to a fair trace $\pi = s_0, s_1, \dots$ of T_A (over Σ_A) s.t. for every $i \in \mathbb{N}$, $s_i|_{\Sigma} = \hat{s}_i$.

Definition 5 (Product system). Given a transition system $T_S = (\Sigma, \iota, \tau)$, a closed FO-LTL formula φ over Σ , a finite set A of FO-LTL formulas over Σ closed under subformulas such that $\neg\varphi \in A$, we define the *product system* of T_S and $\neg\varphi$ over A as the first-order transition system $T_P = (\Sigma_P, \iota_P, \tau_P)$ given by $\Sigma_P = \Sigma_A$, $\iota_P = \iota \wedge FO[\neg\varphi]$ and $\tau_P = \tau \wedge \tau_A$, where $T_A = (\Sigma_A, true, \tau_A)$ is the tableau for A .

Theorem 1. Let T_P be the product system of T_S and $\neg\varphi$ over A as defined in Def. 5. Then $T_S \models \varphi$ iff T_P has no fair traces.

Intuitively, the product system augments the states of T_S with temporal formulas from A , splitting each state into many (often infinitely many) states according to the future behavior of its outgoing traces. Note that Thm. 1 holds already when $A = sub(\neg\varphi)$. However, as we will see, taking a larger set A is useful for proving fair termination via the liveness-to-safety transformation.

4. Liveness-to-Safety with Temporal Prophecy

In this section we present our liveness proof approach using temporal prophecy and a liveness-to-safety transformation. As in earlier approaches, our transformation (i) uses a tableau construction to construct a product transition system equipped with fairness constraints such that the latter has no fair traces iff the temporal property holds of the original system, and (ii) defines a safety property over the product transition system such that safety implies that no fair traces exist (note that the opposite direction does not hold).

The gist of our liveness-to-safety transformation is that we augment the construction of the product transition system with two forms of prophecy detailed in § 4.2. We then use the definition of the safety property from [30]. In the sequel, we first present the safety property and then present the augmentation with temporal prophecy, whose goal is to “refine” the product system such that it will be safe.

4.1. Safety Property: Absence of Abstract Lassos

Given a transition system $T_W = (\Sigma_W, \iota_W, \tau_W)$ with $\Sigma_W \supseteq \Sigma_A$ (e.g., the product system from Def. 5), we define a notion of an abstract lasso, whose absence is a safety property that implies that T_W has no A -fair traces. This section recapitulates material from [30].

The definition of an abstract lasso is based on a dynamic abstraction that is fixed at some point along the trace, henceforth called the *freeze point*. The abstraction function is defined by projecting a state (a first-order structure) into a finite subset of its domain. This finite subset is defined by the union of the *footprints* of all states encountered until the freeze point, where the footprint of a state includes the interpretation it gives all constants from Σ_W . Intuitively, the footprint includes all elements “exposed” in the state, including those “touched” by outgoing transitions.

Definition 6 (Footprint). For a state $s = (\mathcal{D}, \mathcal{I})$ over Σ_W , we define the footprint of s as $f(s) = \{\mathcal{I}(c) \mid c \in \Sigma_W\}$.

For a sequence of states $\pi = s_0, s_1, \dots$ over Σ_W , and an index $i < |\pi|$, we define the footprint of s_0, \dots, s_i as $f(s_0, \dots, s_i) = \bigcup_{j=0}^i f(s_j)$.

Importantly, the footprint of a finite trace is always finite. As a result, an abstraction function that maps each state to the result of projecting it to the footprint of the trace until the freeze point has a finite range.

Definition 7 (Fair Segment). Let $\pi = s_0, s_1, \dots$ be a sequence of states over Σ_W . For $0 \leq i \leq j < |\pi|$, we say the segment $[i, j]$ is fair if for every formula $\Box\psi(\bar{x}) \in A$, and for every assignment σ where every variable is assigned to an element of $f(s_0, \dots, s_i)$, there exists $i \leq k \leq j$ s.t. $s_k, \sigma \models \text{FO}([\Box\psi(\bar{x})] \vee \neg\psi(\bar{x}))$.

Definition 8 (Abstract Lasso). A finite trace s_0, \dots, s_n of T_W is an *abstract lasso* if there are $0 \leq i \leq j < k \leq n$ s.t. the segments $[0, i]$ and $[j, k]$ are fair, and $s_j|_{f(s_0, \dots, s_i)} = s_k|_{f(s_0, \dots, s_i)}$.

Intuitively, in the above definition, i is the *freeze point*, where the abstraction is fixed. The states s_j and s_k are the “repeating states” – states that are indistinguishable by the abstraction that projects them to the footprint $f(s_0, \dots, s_i)$. The segment between j and k , respectively, the segment between 0 and i , meet all the fairness constraints restricted to elements in $f(s_0, \dots, s_j)$, respectively, in $f(s_0)$. Fairness of the segment $[0, i]$ is needed to prevent the freeze point from being chosen too early, thus creating spurious abstract lassos. Note that the absence of abstract lassos is a safety property.

Lemma 3. If T_W has no abstract lassos then it also has no fair traces.

Proof: Assume to the contrary that T_W has a fair trace $\pi = s_0, s_1, \dots$. Let i be the first index such that $[0, i]$ is fair (such an index must exist since the set $f(s_0)$, which determines the relevant fairness constraints is finite). Since $f(s_0, \dots, s_i)$ is also finite, there must exist an infinite subsequence π' of π such that for every s, s' in this subsequence $s|_{f(s_0, \dots, s_i)} = s'|_{f(s_0, \dots, s_i)}$. Let $j \geq i$ be the index in π of

the first state in π' . $f(s_0, \dots, s_j)$ is also finite, hence there exists $k' > j$ such that the segment $[j, k']$ of π is fair. Take k to be the index in π of the first state of $\pi^{k'}$ that is also in π' . Since π' is infinite, such a k must exist. Since $k \geq k'$, the segment $[j, k]$ is also fair. This defines an abstract lasso $s_0, \dots, s_i, \dots, s_j, \dots, s_k$, in contradiction. \square

4.2. Augmenting the Transition System with Temporal Prophecy

In this section we explain how our liveness-to-safety transformation constructs $T_W = (\Sigma_W, \iota_W, \tau_W)$, to which we apply the safety property of § 4.1. Our construction exploits both *temporal prophecy formulas* and *prophecy witnesses*, explained below. For the rest of this section we fix a first-order transition system $T_S = (\Sigma, \iota, \tau)$ and a closed FO-LTL formula φ over Σ that we wish to verify in T_S .

Temporal Prophecy Formulas. First, given a set A of (not necessarily closed) FO-LTL formulas closed under subformula that contains $\neg\varphi$, we construct the product system $T_P = (\Sigma_P, \iota_P, \tau_P)$ defined in Def. 5. By Thm. 1, $T_S \models \varphi$ iff T_P has no fair traces. Note that classical tableau constructions are defined with $A = \text{sub}(\neg\varphi)$, and we allow A to include more formulas. These formulas act as “temporal prophecy variables” in the sense that they split the states of T_S , according to the future behavior of outgoing traces.

While the liveness-to-safety transformation is already sound with $A = \text{sub}(\neg\varphi)$, one of the chief observations of this work is that temporal prophecy formulas improve its precision. These additional formulas in A split the states of T_S into more states in T_P , and they cause some non-determinism of the future trace to be “pulled backwards” (the outgoing traces contain less non-determinism). For example, if $r_{\Box\varphi}$ holds for some elements in the current state, then φ must continue to hold for these elements in the future of the trace. Similarly, for elements where $r_{\Box\varphi}$ does not hold, there will be some time in the future of the trace where φ would not hold for them.

This is exploited by the liveness-to-safety transformation in three ways, eliminating spurious abstract lassos. First, having more temporal formulas in A refines the definition of a fair segment, and postpones the freeze point, thus making the abstraction defined by the footprint up to the freeze point more precise. For example, if $r_{\Box\varphi}$ does not hold for a ground formula φ in the initial state, then the freeze point would be postponed until after φ does not hold for the first time. Second, it strengthens the requirement on the looping segment $s_j \dots s_k$, in a similar way. Third, the additional relations in $\Sigma_P (= \Sigma_A)$ are part of the state as considered by the transformation, and a difference in these relations (projected to the footprint up to the freeze point) is a valid difference. These three ways all played a role in the examples considered in our evaluation.

Prophecy Witnesses. The notion of an abstract lasso, used to define the safety property, considers a finite abstraction according to the footprint, which depends on the constants of the vocabulary. To increase the precision of the

abstraction, we augment the vocabulary with fresh constants that serve as *prophecy witnesses* for existential properties.

To illustrate the idea, consider the formula $\psi(x) = \diamond\Box p(x)$ where x is a free variable. If ψ holds for some element, it is useful to include in the vocabulary a constant that serves as a witness for $\psi(x)$, and whose interpretation will be taken into account by the abstraction. If ψ holds for some x , the interpretation of the constant will be taken from such an x . Otherwise, this constant will be allowed to take any value.

Temporal prophecy witnesses not only refine the abstraction, they can also be used in the inductive invariant. In particular, as demonstrated in the TLB Shutdown example (see § 6), in some cases this allows to avoid quantifier alternation cycles in the verification conditions, leading to decidability of VC checking.

Formally, given a set $B \subseteq A$, we construct $T_W = (\Sigma_W, \iota_W, \tau_W)$ as follows. We extend Σ_P to Σ_W by adding fresh constant symbols c_1, \dots, c_n for every formula $\psi(x_1, \dots, x_n) \in B$. We denote by C the set of new constants, i.e., $\Sigma_W = \Sigma_P \cup C$. The transition relation formula is extended to keep the new constants unchanged, i.e. $\tau_W = \tau_P \wedge \bigwedge_{c \in C} c = c'$, and we define ι_W by

$$\iota_W = \iota_P \wedge \text{FO}[(\exists x_1, \dots, x_n. \psi(x_1, \dots, x_n)) \rightarrow \psi(c_1, \dots, c_n)]$$

Namely, c_1, \dots, c_n are required to serve as witnesses for $\psi(x_1, \dots, x_n)$ in case it holds in the initial state for some elements, and otherwise they may get any interpretation at the initial state, after which their interpretation remains unchanged. Adding these fresh constants and their defining formulas to the initial state is a conservative extension, in the sense that every fair trace of T_P can be extended to a fair trace of T_W (fairness of traces over $\Sigma_W \supseteq \Sigma_A$ is defined as in Def. 4), and every fair trace of T_W can be projected to a fair trace of T_P . As such we have the following:

Lemma 4. Let $T_P = (\Sigma_P, \iota_P, \tau_P)$ and $T_W = (\Sigma_W, \iota_W, \tau_W)$ be defined as above. Then T_P has no fair traces iff T_W has no fair traces.

The overall soundness of the liveness-to-safety transformation is given by the following theorem.

Theorem 2 (Soundness). Given a first-order transition system T_S and a closed FO-LTL formula φ both over Σ , and given a set of temporal prophecy formulas A over Σ that contains $\neg\varphi$ and is closed under subformula, and a set of temporal prophecy witness formulas $B \subseteq A$, if T_W as defined above does not contain an abstract lasso, then $T_S \models \varphi$.

4.3. The Ticket Example

In this section we show in greater detail how prophecy increases the power of the liveness-to-safety transformation. As an illustration we return to the ticket example of Fig. 1. As explained in § 2, in this example the liveness-to-safety transformation without temporal prophecy fails (similarly to [30, §5.2]), but it succeeds when adding suitable temporal prophecy.

To model the ticket example as a first-order transition system, we use a vocabulary with two sorts: *thread* and *number*. The first represents threads, and the second represents ticket values and counter values. The vocabulary also includes a static binary relation symbol \leq : *number, number*, with suitable first-order axioms to make it a total order. (for more details about modeling systems in first-order logic see e.g. [31].) The state of the system is modeled by unary relations for the program counter: *idle, wait, critical*, constant symbols of sort *number* for the global variables n, s , and functions from *thread* to *number* for the local variables m, c . The vocabulary also includes a unary relation *scheduled*, which holds the last scheduled thread.

Next we show that when adding the temporal prophecy formula $\exists x. \diamond\Box \text{critical}(x)$ to the tableau construction, no abstract lasso exists in the augmented transition system, hence the liveness-to-safety transformation succeeds to prove the property. Formally, in this case, A includes the following two formulas and their subformulas:

$$\neg((\exists x. \neg\Box\neg\Box\text{-scheduled}(x)) \vee \neg\exists x. \neg\Box(\neg\text{wait}(x) \vee \neg\Box\neg\text{critical}(x))) \\ \exists x. \neg\Box\neg\Box\text{critical}(x)$$

And $B = \{\neg\Box(\neg\text{wait}(x) \vee \neg\Box\neg\text{critical}(x)), \neg\Box\neg\Box\text{critical}(x)\}$. Therefore, Σ_W extends the original vocabulary with the following 6 unary relations:

$$r_{\Box\neg\text{scheduled}(x)}, r_{\Box\neg\Box\neg\text{scheduled}(x)}, r_{\Box\neg\text{critical}(x)}, \\ r_{\Box\neg\text{wait}(x) \vee \neg\Box\neg\text{critical}(x)}, r_{\Box\text{critical}(x)}, r_{\Box\neg\Box\text{critical}(x)}$$

as well as two constants for prophecy witnesses: c_1 for $\neg\Box(\neg\text{wait}(x) \vee \neg\Box\neg\text{critical}(x))$, and c_2 for $\neg\Box\neg\Box\text{critical}(x)$.

We now explain why there is no abstract lasso. To do this, we show that the tableau construction, combined with the dynamic abstraction and the fair segment requirements, result in the same reasoning that was presented informally in § 2.

First, observe that from the definition of c_1 and the negation of the liveness property (both assumed by ι_W), we have that the initial state $s_0 \models \text{FO}[\neg\Box(\neg\text{wait}(c_1) \vee \neg\Box\neg\text{critical}(c_1))]$. For brevity, denote $p = (\neg\text{wait}(c_1) \vee \neg\Box\neg\text{critical}(c_1))$, so we have $s_0 \models \text{FO}[\neg\Box p]$, i.e., $s_0 \models \neg r_{\Box p}$. Since c_1 is also in the footprint of the initial state, the fair segment requirement ensures that the freeze point can only happen after encountering a state satisfying: $\text{FO}[(\Box p) \vee \neg p] \equiv r_{\Box p} \vee \text{FO}[\neg p]$. Recall that the transition relation of the tableau (τ_A), ensures $(r_{\Box p}) \leftrightarrow (\text{FO}[p] \wedge r_{\Box p}')$. Therefore, on update from a state satisfying $\neg r_{\Box p}$ to a state satisfying $r_{\Box p}$ can only happen if the pre-state satisfies $\text{FO}[\neg p]$. Therefore, the freeze point must come after encountering a state that satisfies $\text{FO}[\neg p] \equiv \text{wait}(c_1) \wedge r_{\Box\neg\text{critical}(c_1)}$. From the freeze point onward, τ_A will ensure both $r_{\Box\neg\text{critical}(c_1)}$ and $\neg\text{critical}(c_1)$ continue to hold, so c_1 will stay in *wait* (since the protocol does not allow to go from *wait* to anything but *critical*). So, we see that the mechanism of the tableau, combined with the temporal prophecy witness and the fair segment requirement, ensures that the freeze point happens after c_1

makes a request that is never granted. This will ensure that the footprint used for the dynamic abstraction will include all threads ahead of c_1 in line, i.e., those with smaller ticket numbers.

As for c_2 , the initial state will either satisfy $\text{FO}[\neg\Box\neg\Box\text{critical}(c_2)]$ or it would satisfy $\text{FO}[\neg\exists x.\neg\Box\neg\Box\text{critical}(x)]$. In the first case, by an argument similar to the one used above for c_1 , the freeze point will happen after c_2 enters the critical section and then stays in it. Therefore, the footprint used for the dynamic abstraction will include all numbers smaller than q of c_2 when it enters the critical section¹. Since c_2 is required to be scheduled between the repeating states (again by the tableau construction and the fair segment requirement), its value for q will be decreased, and this will be visible in the dynamic abstraction. Thus, in this case, an abstract lasso is not possible.

In the second case the initial state satisfies $\text{FO}[\neg\exists x.\neg\Box\neg\Box\text{critical}(x)]$. By a similar argument that combines the tableau with the fair segment requirement for the repeating states, we will obtain that between the repeating states, any thread in the footprint of the first repeating state, must both be scheduled and visit a state outside the critical section. In particular, this includes all threads that are ahead of c_1 in line. This entails a change to the program counter of one of them (the one that had a ticket number equal to the service number at the first repeating state), which will be visible in the abstraction. Thus, an abstract lasso is not possible in this case either.

5. Closure Under First-Order Reasoning

The transformation from temporal verification to safety verification developed in § 4 introduces an abstraction, and incurs a loss of precision. That is, for some systems and properties, liveness holds but the safety of the resulting system does not hold, no matter what temporal prophecy is used. (This is unavoidable for a transformation from arbitrary FO-LTL properties to safety properties [30].) However, in this section, we show that the set of instances for which the transformation can be made precise (via temporal prophecy) is closed under first-order reasoning. This is unlike the transformation of [30]. It shows that the use of temporal prophecy results in a particular kind of robustness.

We consider a proof system in which the above transformation is performed and the resulting safety property is checked by an oracle. That is, for a transition system T_S and a temporal property φ (a closed FO-LTL formula), we write $T_S \vdash \varphi$ if there exist finite sets of FO-LTL formulas A and B satisfying the conditions of Thm. 2, such that resulting transition system T_W is safe, i.e., does not contain an abstract lasso. We now show that the relation \vdash satisfies a powerful closure property.

1. When modeling natural numbers in first-order logic, the footprint is adjusted to include all numbers lower than any constant (still being a finite set).

Theorem 3 (Closure under first-order reasoning). Let T_S be a transition system, and $\psi, \varphi_1, \dots, \varphi_n$ be closed FO-LTL formulas, such that $\text{FO}[\varphi_1 \wedge \dots \wedge \varphi_n] \models \text{FO}[\psi]$. If $T_S \vdash \varphi_i$ for all $1 \leq i \leq n$, then $T_S \vdash \psi$.

The condition that $\text{FO}[\varphi_1 \wedge \dots \wedge \varphi_n] \models \text{FO}[\psi]$ means that $\varphi_1 \wedge \dots \wedge \varphi_n$ entails ψ when using only first-order reasoning, and treating temporal operators as uninterpreted. The theorem states that provability using the liveness-to-safety transformation is closed under such reasoning. Two special cases of Thm. 3 given by the following corollaries:

Corollary 1 (Modus Ponens). If T_S is a transition system and φ and ψ are closed FO-LTL formulas such that $T_S \vdash \varphi$ and $T_S \vdash \varphi \rightarrow \psi$, then $T_S \vdash \psi$.

Corollary 2 (Cut). If T_S is a transition system and φ and ψ are closed FO-LTL formulas such that $T_S \vdash \varphi \rightarrow \psi$ and $T_S \vdash \neg\varphi \rightarrow \psi$, then $T_S \vdash \psi$.

Proof of Thm. 3: In the proof we use the notation $T_W(T_S, \varphi, A, B)$ to denote the transition system constructed for T_S and φ when using A, B as temporal prophecy formulas. Likewise, we refer to the vocabulary, initial states and transition relation formulas of the transition system as $\Sigma_W(T_S, \varphi, A, B)$, $\iota_W(T_S, \varphi, A, B)$, and $\tau_W(T_S, \varphi, A, B)$, respectively. Let $(A_1, B_1), \dots, (A_n, B_n)$ be such that $T_W(T_S, \varphi_i, A_i, B_i)$ has no abstract lasso, for every $1 \leq i \leq n$. Now, let $A = \bigcup_{i=1}^n A_i$ and $B = \bigcup_{i=1}^n B_i$. We show that $T_W(T_S, \psi, A, B)$ has no abstract lasso. Assume to the contrary that $s_0, \dots, s_i, \dots, s_j, \dots, s_k, \dots, s_n$ is an abstract lasso for $T_W(T_S, \psi, A, B)$. Since $s_0 \models \neg\text{FO}[\psi]$, and since $\text{FO}[\varphi_1 \wedge \dots \wedge \varphi_n] \models \text{FO}[\psi]$, there must be some $1 \leq \ell \leq n$ s.t. $s_0 \models \neg\text{FO}[\varphi_\ell]$. Denote $\Sigma' = \Sigma_W(T_S, \varphi_\ell, A_\ell, B_\ell)$. Now, $s_0|_{\Sigma'}, \dots, s_i|_{\Sigma'}, \dots, s_j|_{\Sigma'}, \dots, s_k|_{\Sigma'}, \dots, s_n|_{\Sigma'}$ is an abstract lasso of $T_W(T_S, \varphi_\ell, A_\ell, B_\ell)$, which is a contradiction. To see that, we first simplify the notation and denote $s_m|_{\Sigma'}$ by \hat{s}_m . The footprint $f(s_0, \dots, s_i)$ contains more elements than the footprint $f(\hat{s}_0, \dots, \hat{s}_i)$, since $\Sigma_W(T_S, \psi, A, B) \supseteq \Sigma_W(T_S, \varphi_\ell, A_\ell, B_\ell)$. Therefore, given that $s_j|_{f(s_0, \dots, s_i)} = s_k|_{f(s_0, \dots, s_i)}$, we have that $\hat{s}_j|_{f(\hat{s}_0, \dots, \hat{s}_i)} = \hat{s}_k|_{f(\hat{s}_0, \dots, \hat{s}_i)}$ as well. Moreover, the fairness constraints in $T_W(T_S, \varphi_\ell, A_\ell, B_\ell)$, determined by A_ℓ , are a subset of those in $T_W(T_S, \psi, A, B)$, determined by A , so the segments $[0, i]$ and $[j, k]$ are also fair in $T_W(T_S, \varphi_\ell, A_\ell, B_\ell)$. \square

The proof of Thm. 3 sheds more light on the power of using temporal prophecy formulas that are not subformulas of the temporal property to prove. In particular, the theorem does not hold if A is restricted to subformulas of the temporal proof goal.

6. Implementation & Evaluation

We have implemented our approach for temporal verification and integrated it into the Ivy deductive verification system [31]. This allows the user to model the transition system in the Ivy language (which internally translates into

a first-order transition system), and express temporal properties directly in FO-LTL. In our implementation, the safety property that results from the liveness-to-safety transformation is proven by a suitable inductive invariant, provided by the user. To facilitate this process, Ivy internally constructs a suitable monitor for the safety property, i.e., the absence of abstract lasso’s in T_W . The user then provides an inductive invariant for T_W composed with this monitor. The monitor keeps track of the footprint and the fairness constraints, and non-deterministically selects the freeze point and repeated states of the abstract lasso. Similar to the construction of [5], the monitor keeps a shadow copy of the “saved state”, which is the first of the two repeated states. These are maintained via designated relation symbols (in addition to Σ_W). The user’s inductive invariant must then prove that it is impossible for the monitor to detect an abstract lasso.

Mining Temporal Prophecy from the Invariant. As presented in previous sections, our liveness-to-safety transformation is parameterized by sets of formulas A and B . In the implementation, these sets are implicit, and are extracted automatically from the inductive invariant provided by the user. Namely, the inductive invariant provided by the user contains temporal formulas, and also prophecy witness constants, where every temporal formula $\Box\varphi$ is a shorthand (and is internally rewritten to) $r\Box\varphi$. The set A to be used in the construction is defined by all the temporal subformulas that appear in the inductive invariant (and all their subformulas), and the set B is defined according to the prophecy witness constants that are used in the inductive invariant.

In particular, the user’s invariant may refer to the satisfaction of each fairness constraint $\text{FO}[\Box\varphi \vee \neg\varphi]$ for $\Box\varphi \in A$, both before the freeze point and between the repeated states, via a convenient syntax provided by Ivy.

Interacting with Ivy. If the user provides an inductive invariant that is not inductive, Ivy presents a graphical counterexample to induction. This guides the user to adjust the inductive invariant, which may also lead to new formulas being added to A or B , if the user adds new temporal formulas or prophecy witnesses to the inductive invariant. In this process, the user’s mental image is of a liveness-to-safety transformation where A and B include all (countably many) FO-LTL formulas over the system’s vocabulary, so the user is free to use any temporal formula, or prophecy witness for any formula. However, since the user’s inductive invariant is a finite formula, the liveness-to-safety transformation needs only to be applied to finite A and B , and the infinite A and B are just a mental model.

We have used our implementation to prove liveness for several challenging examples, summarized in Fig. 2. We focused on examples that were beyond reach for the liveness-to-safety transformation of [30]. In [30], such examples were handled using a nesting structure. Our experience shows that with temporal prophecy, the invariants are simpler than with a nesting structure (for additional comparison with nesting structure see § 7). For all examples we considered, the verification conditions are in a decidable fragment of first-order logic which is supported by Z3 (the stratified

Protocol	# A	# B	# LOC	# C	FO-LTL	t [sec]
Ticket w/ Task Queues	1	2	90	60	22%	9.4
Alternating Bit Protocol	4	1	143	70	40%	32
TLB Shutdown	6	3	468	102	49%	283

Figure 2. Protocols for which we verified liveness. For each protocol, # **A** reports the number of temporal prophecy formulas used. # **B** reports the number of prophecy witnesses used. # **LOC** reports the number of lines of code for the system model (without proof) in Ivy’s modeling language. # **C** reports the number of conjectures used in the inductive invariant (a typical conjecture is one or few lines). **FO-LTL** reports the fraction of the conjectures that use temporal formulas. Finally, **t** reports the run time (in seconds) for checking the verification conditions using Ivy and Z3. The experiments were performed on a laptop running 64-bit Linux, with a Core-i7 1.8 GHz CPU, using Z3 version 4.6.0.

extension of EPR [19], [31]). Interestingly, for the TLB shutdown example, the proof presented in [30] (using a nesting structure) required non-stratified quantifier alternation, which is eliminated by the use of temporal prophecy witnesses. Due to the decidability of verification conditions, Z3 behaves predictably, and whenever the invariant is not inductive it produces a finite counterexample to induction, which Ivy presents graphically. Our experience shows that the graphical counterexamples provide valuable guidance towards finding an inductive invariant, and also for coming up with temporal prophecy formulas as needed. Below we provide more detail on each example.

Ticket. The ticket example has been discussed in § 1, and § 4.3 contains more details about its proof with temporal prophecy, using a single temporal prophecy formula and two prophecy witness constants. To give a flavor of what the proof looks like in Ivy, we present a couple of the conjectures that make up the inductive invariant for the resulting system, in Ivy’s syntax. In Ivy, the prefix `l2s` indicates symbols that are introduced by the liveness-to-safety transformation. Some conjectures are needed to state that the footprint used in the dynamic abstraction contains enough elements. An example of such a conjecture is:

```
l2s_frozen & (globally critical(c2)) ->
forall N. N <= q(c2) -> l2s_a(N)
```

This conjecture states that after the freeze point (indicated by the special symbol `l2s_frozen`), if the prophecy witness `c2` (which is the prophecy witness defined for $\Diamond\Box\text{critical}(x)$) is globally in the critical section, then the finite domain of the frozen abstraction (stored in the unary relation `l2s_a`) contains all numbers up the `c2`’s value for `q`. Other conjectures are needed to show that the current state is different from the saved state. One example is:

```
l2s_saved & (globally critical(c2)) &
~($l2s_w X. scheduled(X))(c2) ->
q(c2) ~= ($l2s_s X. q(X))(c2)
```

The special operator `$l2s_w` lets the user query whether a fairness constraint has been encountered, and `$l2s_s` exposes to the user the saved state (both syntactically λ -like binders). This conjecture states that after we saved a shadow state (indicated by `l2s_save`), if the prophecy witness `c2` is globally in the critical section, and if we

have encountered the fairness constraints associated with $scheduled(x) \vee \Box \neg scheduled(x)$ instantiated for c_2 (which can only happen after c_2 has been scheduled), then the current value c_2 has for q is different from the same value in the shadow state.

Alternating Bit Protocol. The alternating bit protocol is a classic communication algorithm for transition of messages using lossy first-in-first-out (FIFO) channels. The protocol uses two channels: a data channel from the sender to the receiver, and an acknowledgment channel from the receiver to the sender. The sender and the receiver each have a state bit, and messages include a bit that functions as a “sequence number”. We assume that the sender has an (infinite) array of values to send, which is filled by some independent process. The liveness property we wish to prove is that every value entered into the sender array is eventually received by the receiver.

The protocol is live under fair scheduling assumptions, as well as standard fairness constraints for the channels: if messages are infinitely often sent, then messages are infinitely often received. This makes the structure of the temporal property more involved. Formally, the liveness property we prove is:

$$\begin{aligned} & (\Box \diamond sender_scheduled) \wedge (\Box \diamond receiver_scheduled) \wedge \\ & ((\Box \diamond data_sent) \rightarrow (\Box \diamond data_received)) \wedge \\ & ((\Box \diamond ack_sent) \rightarrow (\Box \diamond ack_received)) \rightarrow \\ & \forall x. \Box (sender_array(x) \neq \perp \rightarrow \diamond receiver_array(x) \neq \perp) \end{aligned}$$

This property cannot be proven without temporal prophecy. However, it can be proven using 4 temporal prophecy formulas: $\{\diamond \Box (sender_bit = s \wedge receiver_bit = r) \mid s, r \in \{0, 1\}\}$. Intuitively, these formulas make a distinction between traces in which the sender and receiver bits eventually become fixed, and traces in which they change infinitely often.

TLB Shootdown. The TLB shutdown algorithm [6] is used (e.g. in the Mach operating system) to maintain consistency of Translation Look-aside Buffers (TLB) across processors. When some processor (dubbed the initiator) changes the page table, it interrupts all other processors currently using the page table (dubbed the responders) and waits for them to receive the interrupt before making changes. The liveness property we prove is that no processor can become stuck either as an initiator or as a responder (formally, it will respond or initiate infinitely often). This liveness depends on fair scheduling assumptions, as well as strong fairness assumptions for the page table locks used by the protocol. We use one witness for the process that does not satisfy the liveness property. Another witness is used for a pagemap that is never unlocked, if this exists. A third witness is used for a process that possibly gets stuck while holding the lock blocking the first process. We use six prophecy formulas to case split on when some process may get stuck. Two of them are used for the two loops in the initiator to distinguish the cases whether the process that hogs the lock gets stuck there. They are of the form $\diamond \Box pc(c_2) \in \{i_3, \dots, i_8\}$. Two are used for the two lock instructions to indicate that the first process gets stuck: $\diamond \Box pc(c_1) = i_2$. And two are used

for the second and third witness to indicate whether such a witness exists, e.g., $\diamond \Box plock(c_3)$. Compared to the proof of [30], our proof is simpler due to the temporal prophecy, and avoids non-stratified quantifier alternation, resulting in decidable verification conditions.

7. Related Work

Prophecy variables were first introduced in [2], in the context of refinement mappings. There, prophecy variables are required to range over a finite domain to ensure soundness. Our notion of prophecy via first-order temporal formulas and witness constants does not meet this criterion, but is still sound as assured by Thm. 2. In [25], LTL formulas are used to define prophecy variables in a way that is similar to ours, but only to show refinement between finite-state processes. We use temporal prophecy defined by FO-LTL formulas in the context of infinite-state systems. Furthermore, we consider a liveness-to-safety transformation (rather than refinement mappings), which can be seen as a proof system for FO-LTL.

The liveness-to-safety transformation based on dynamic abstraction, but *without temporal prophecy*, was introduced in [30]. There, a *nesting structure* was used to increase the power of the transformation. A nesting structure is defined by the user (via first-order formulas), and has the effect of splitting the transition system into levels (analogous to nested loops) and proving each level separately. Temporal prophecy as we introduce here is more general, and in particular, any proof that is possible with a nesting structure, is also possible with temporal prophecy (by adding a temporal prophecy formula $\diamond \Box \delta$ for every nesting level, defined by δ). Moreover, the nesting structure does not admit cut elimination or closure under first-order reasoning, and is therefore less robust.

One effect of prophecy is to split cases in the proof on some aspect of the future. This very general idea occurs in various approaches to liveness, particularly in the large body of work on lexicographic or disjunctive rankings for termination [4], [7], [8], [11], [12], [14], [18], [20], [21], [23], [26], [27], [33], [34], [35], [36], [37], [38]. In the work of [22], the partitioning of the space of potentially infinite executions is based on the *a priori* decomposition of regular expressions for iterated loop segments. Often the partitioning here amounts to a split according to a fairness condition (“command a is taken infinitely often or it is not”). The partitioning is constructed dynamically (and represented explicitly through a union of Buchi automata) in [24] (for termination), in [15] (for liveness), and in [17] (for liveness of parameterized systems). None of these works uses a temporal tableau construction to partition the space of futures, however.

Here, we use prophecy to, in effect, partially determinize a system by making non-deterministic choices earlier in an execution. This same effect was used for a different purpose in refining an abstraction from LTL to ACTL [10] and checking CTL* properties [9]. The prophecy in this case relates only to the next transition and is not expressed

temporally. The method of “temporal case splitting” in [29] can also be seen as a way to introduce prophecy variables to increase the precision of an abstraction, though in that case the transformation was to finite-state liveness, not infinite-state safety. Moreover, it only introduces temporal witnesses.

We have considered only proof methods that transform liveness to safety (which includes the classical ranking approach for while loops). There are approaches, however, which do *not* transform liveness to safety. For example, the approaches in [3], [14], [39] are essentially forms of widening in a CTL-style backwards fixpoint iteration. It is not clear to what extent temporal prophecy might be useful in increasing the precision of such abstractions, but it may be an interesting topic for future research.

8. Conclusion

We have seen that the addition of prophecy variables in the form of temporal formulas can increase the precision of liveness-to-safety transformations for infinite-state systems. The prophecy variables are derived from additional temporal formulas that in our implementation were mined from the invariants a user provides to prove the safety property. This approach is effective for proving challenging examples. By increasing the precision of the dynamic abstraction, it avoided the need to decompose the proof into nested termination arguments, reducing the human effort of proof construction. Though completeness is not possible, we saw that the additional expressiveness of temporal prophecy provides a cut elimination property. While we considered temporal prophecy using a particular liveness-to-safety construction (based on dynamic abstraction), it seems reasonable to expect that the tableau-based approach would apply to other constructions and abstractions, including constructions based on rankings and well-founded relations. Because our approach relies on an inductive invariant supplied by the user, it requires the user to understand the liveness-to-safety transformation and it requires both cleverness and a deep understanding of the protocol. For this reason, a possible avenue for future research would be to explore invariant synthesis techniques, and in particular ones that account for refinement due to temporal prophecy.

Acknowledgements. We thank the anonymous referees for insightful comments which improved this paper. Padon was supported by Google under a PhD fellowship. Padon and Sagiv were supported by the European Research Council under the European Union’s Seventh Framework Program (FP7/2007–2013) / ERC grant agreement no. [321174-VSSC]. This publication is part of a project that has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No [759102-SVIS]). The research was partially supported by Len Blavatnik and the Blavatnik Family foundation, and by the Blavatnik Interdisciplinary Cyber Research Center, Tel Aviv University. This material is based upon work supported by the United States-Israel Binational Science Foundation (BSF) grants No. 2016260 and 2012259.

References

- [1] M. Abadi, “The power of temporal proofs,” *Theor. Comput. Sci.*, vol. 65, no. 1, pp. 35–83, 1989. [Online]. Available: [https://doi.org/10.1016/0304-3975\(89\)90138-2](https://doi.org/10.1016/0304-3975(89)90138-2)
- [2] M. Abadi and L. Lamport, “The existence of refinement mappings,” *Theor. Comput. Sci.*, vol. 82, no. 2, pp. 253–284, 1991. [Online]. Available: [https://doi.org/10.1016/0304-3975\(91\)90224-P](https://doi.org/10.1016/0304-3975(91)90224-P)
- [3] P. A. Abdulla, B. Jonsson, A. Rezine, and M. Saksena, “Proving liveness by backwards reachability,” in *CONCUR*, ser. Lecture Notes in Computer Science, vol. 4137. Springer, 2006, pp. 95–109.
- [4] D. Babic, A. J. Hu, Z. Rakamaric, and B. Cook, “Proving termination by divergence,” in *SEFM*, 2007, pp. 93–102.
- [5] A. Biere, C. Artho, and V. Schuppan, “Liveness checking as safety checking,” *Electr. Notes Theor. Comput. Sci.*, vol. 66, no. 2, pp. 160–177, 2002.
- [6] D. L. Black, R. F. Rashid, D. B. Golub, and C. R. Hill, “Translation lookaside buffer consistency: A software approach,” in *Proceedings of the Third International Conference on Architectural Support for Programming Languages and Operating Systems*, ser. ASPLOS III. New York, NY, USA: ACM, 1989, pp. 113–122. [Online]. Available: <http://doi.acm.org/10.1145/70082.68193>
- [7] M. Brockschmidt, B. Cook, and C. Fuhs, “Better termination proving through cooperation,” in *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, 2013, pp. 413–429.
- [8] B. Cook, A. Podelski, and A. Rybalchenko, “Termination proofs for systems code,” in *PLDI*, 2006, pp. 415–426.
- [9] B. Cook, H. Khlaaf, and N. Piterman, “On automation of cti* verification for infinite-state systems,” in *Computer Aided Verification - 27th International Conference, CAV 2015, San Francisco, CA, USA, July 18-24, 2015, Proceedings, Part I*, ser. Lecture Notes in Computer Science, D. Kroening and C. S. Pasareanu, Eds., vol. 9206. Springer, 2015, pp. 13–29. [Online]. Available: https://doi.org/10.1007/978-3-319-21690-4_2
- [10] B. Cook and E. Koskinen, “Making prophecies with decision predicates,” in *Proceedings of the 38th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2011, Austin, TX, USA, January 26-28, 2011*, T. Ball and M. Sagiv, Eds. ACM, 2011, pp. 399–410. [Online]. Available: <http://doi.acm.org/10.1145/1926385.1926431>
- [11] B. Cook, A. Podelski, and A. Rybalchenko, “Proving program termination,” *Commun. ACM*, vol. 54, no. 5, pp. 88–98, 2011.
- [12] B. Cook, A. See, and F. Zuleger, “Ramsey vs. lexicographic termination proving,” in *TACAS*, 2013, pp. 47–61.
- [13] J. Corbet, “Ticket spinlocks,” <https://lwn.net/Articles/267968/>, 2008.
- [14] P. Cousot and R. Cousot, “An abstract interpretation framework for termination,” in *POPL*, 2012, pp. 245–258.
- [15] D. Dietsch, M. Heizmann, V. Langenfeld, and A. Podelski, “Fairness modulo theory: A new approach to LTL software model checking,” in *CAV*, ser. Lecture Notes in Computer Science, vol. 9206. Springer, 2015, pp. 49–66.
- [16] Y. Fang, K. L. McMillan, A. Pnueli, and L. D. Zuck, “Liveness by invisible invariants,” in *Formal Techniques for Networked and Distributed Systems - FORTE 2006, 26th IFIP WG 6.1 International Conference, Paris, France, September 26-29, 2006.*, ser. Lecture Notes in Computer Science, E. Najm, J. Pradat-Peyre, and V. Donzeau-Gouge, Eds., vol. 4229. Springer, 2006, pp. 356–371. [Online]. Available: https://doi.org/10.1007/11888116_26
- [17] A. Farzan, Z. Kincaid, and A. Podelski, “Proving liveness of parameterized programs,” in *LICS*. ACM, 2016, pp. 185–196.
- [18] P. Ganty and S. Genaim, “Proving termination starting from the end,” in *Computer Aided Verification - 25th International Conference, CAV 2013, Saint Petersburg, Russia, July 13-19, 2013. Proceedings*, 2013, pp. 397–412.

- [19] Y. Ge and L. D. Moura, “Complete instantiation for quantified formulas in satisfiability modulo theories,” in *International Conference on Computer Aided Verification*. Springer, 2009, pp. 306–320.
- [20] J. Giesl, R. Thiemann, P. Schneider-Kamp, and S. Falke, “Automated termination proofs with AProVE,” in *RTA*, 2004, pp. 210–220.
- [21] S. Grebenshchikov, N. P. Lopes, C. Popeea, and A. Rybalchenko, “Synthesizing software verifiers from proof rules,” in *PLDI*, 2012, pp. 405–416.
- [22] S. Gulwani, S. Jain, and E. Koskinen, “Control-flow refinement and progress invariants for bound analysis,” in *PLDI*, 2009, pp. 375–385.
- [23] W. R. Harris, A. Lal, A. V. Nori, and S. K. Rajamani, “Alternation for termination,” in *Static Analysis - 17th International Symposium, SAS 2010, Perpignan, France, September 14-16, 2010. Proceedings*, 2010, pp. 304–319.
- [24] M. Heizmann, J. Hoenicke, and A. Podelski, “Termination analysis by learning terminating programs,” in *CAV*, ser. Lecture Notes in Computer Science, vol. 8559. Springer, 2014, pp. 797–813.
- [25] Y. Kesten, A. Pnueli, E. Shahar, and L. D. Zuck, “Network invariants in action,” in *Proceedings of the 13th International Conference on Concurrency Theory*, ser. CONCUR ’02. Berlin, Heidelberg: Springer-Verlag, 2002, pp. 101–115. [Online]. Available: <http://dl.acm.org/citation.cfm?id=646737.701938>
- [26] D. Kroening, N. Sharygina, A. Tsitovich, and C. M. Wintersteiger, “Termination analysis with compositional transition invariants,” in *Computer Aided Verification, 22nd International Conference, CAV 2010, Edinburgh, UK, July 15-19, 2010. Proceedings*, 2010, pp. 89–103.
- [27] W. Lee, B. Wang, and K. Yi, “Termination analysis with algorithmic learning,” in *Computer Aided Verification - 24th International Conference, CAV 2012, Berkeley, CA, USA, July 7-13, 2012 Proceedings*, 2012, pp. 88–104.
- [28] R. Manevich, B. Dogadov, and N. Rinetzky, “From shape analysis to termination analysis in linear time,” in *CAV (1)*, ser. Lecture Notes in Computer Science, vol. 9779. Springer, 2016, pp. 426–446.
- [29] K. L. McMillan, “A methodology for hardware verification using compositional model checking,” *Sci. Comput. Program.*, vol. 37, no. 1-3, pp. 279–309, 2000. [Online]. Available: [https://doi.org/10.1016/S0167-6423\(99\)00030-1](https://doi.org/10.1016/S0167-6423(99)00030-1)
- [30] O. Padon, J. Hoenicke, G. Losa, A. Podelski, M. Sagiv, and S. Shoham, “Reducing liveness to safety in first-order logic,” *PACMPL*, vol. 2, no. POPL, pp. 26:1–26:33, 2018. [Online]. Available: <http://doi.acm.org/10.1145/3158114>
- [31] O. Padon, K. L. McMillan, A. Panda, M. Sagiv, and S. Shoham, “Ivy: safety verification by interactive generalization,” in *Proceedings of the 37th ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI 2016, Santa Barbara, CA, USA, June 13-17, 2016*, 2016, pp. 614–630.
- [32] A. Pnueli and E. Shahar, “Liveness and acceleration in parameterized verification,” in *CAV*, ser. Lecture Notes in Computer Science, vol. 1855. Springer, 2000, pp. 328–343.
- [33] A. Podelski and A. Rybalchenko, “Transition invariants,” in *19th IEEE Symposium on Logic in Computer Science (LICS 2004)*, 14-17 July 2004, Turku, Finland, *Proceedings*, 2004, pp. 32–41.
- [34] —, “Transition predicate abstraction and fair termination,” in *Proceedings of the 32nd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2005, Long Beach, California, USA, January 12-14, 2005*, 2005, pp. 132–144.
- [35] C. Urban, “The abstract domain of segmented ranking functions,” in *SAS*, ser. Lecture Notes in Computer Science, vol. 7935. Springer, 2013, pp. 43–62.
- [36] C. Urban, A. Gurfinkel, and T. Kahsai, “Synthesizing ranking functions from bits and pieces,” in *TACAS*, ser. Lecture Notes in Computer Science, vol. 9636. Springer, 2016, pp. 54–70.
- [37] C. Urban and A. Miné, “An abstract domain to infer ordinal-valued ranking functions,” in *Programming Languages and Systems - 23rd European Symposium on Programming, ESOP 2014, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2014, Grenoble, France, April 5-13, 2014, Proceedings*, 2014, pp. 412–431.
- [38] —, “A decision tree abstract domain for proving conditional termination,” in *SAS*, ser. Lecture Notes in Computer Science, vol. 8723. Springer, 2014, pp. 302–318.
- [39] —, “Inference of ranking functions for proving temporal properties by abstract interpretation,” *Computer Languages, Systems & Structures*, vol. 47, pp. 77–103, 2017.