

Monitoring Arithmetic Temporal Properties on Finite Traces

Paolo Felli,¹ Marco Montali,² Fabio Patrizi,³ Sarah Winkler²

¹Università di Bologna – Bologna – Italy, paolo.felli@unibo.it

²Free University of Bozen-Bolzano – Bolzano – Italy, {montali,winkler}@inf.unibz.it

³Sapienza Università di Roma – Rome – Italy, patrizi@diag.uniroma1.it

Abstract

We study the monitoring of linear-time arithmetic properties against finite traces generated by an unknown dynamic system. The monitoring state is determined by considering at once the trace prefix seen so far, and all its possible finite-length, future continuations. This makes monitoring at least as hard as satisfiability and validity. Traces consist of finite sequences of assignments of a fixed set of variables to numerical values. Properties are specified in a logic we call $ALTL_f$, combining LTL_f (LTL on finite traces) with linear arithmetic constraints that may carry *lookahead*, i.e., variables may be compared over multiple instants of the trace. While the monitoring problem for this setting is undecidable in general, we show decidability for (a) properties without lookahead, and (b) properties with lookahead that satisfy the abstract, semantic condition of *finite summary*, studied before in the context of model checking. We then single out concrete, practically relevant classes of constraints guaranteeing finite summary. Feasibility is witnessed by a prototype implementation.

Introduction

Dynamic systems in AI are often constituted by autonomous agents and heterogeneous components whose internal specification is either unknown or not accessible, so that well-known techniques to ascertain their correctness at design-time, such as model checking and testing, are not applicable. This calls for approaches to check desired properties at runtime, by *monitoring* the executions of the black-box system under scrutiny. A widely known, solid approach to monitoring is that of *runtime verification* (RV), where given a logical property, the monitor emits a provably correct verdict (Leucker and Schallhart 2009). Formalisms based on LTL and its extensions have indeed been extensively applied when monitoring multi-agent systems (Dastani, Torroni, and Yorke-Smith 2018), software components (Leucker and Schallhart 2009) and business processes (Ly et al. 2015).

Leucker and Schallhart (2009) highlight two essential semantic desiderata of linear-time monitors. First, a monitored trace has a finite length, which calls for a finite-trace semantics. Second, a trace is the prefix of an unknown, full trace. Thus, the verdict of the monitor should be *anticipatory* (Barocchi et al. 2018), i.e., depend not only on the data seen so


far, but also on its (infinitely many) possible continuations. In addition, runtime monitors need to do as little computation as possible during monitoring, and therefore all the required machinery and anticipatory test conditions need to be pre-computed. Similarly, the monitor itself cannot depend on the specific trace being monitored, but should allow for monitoring multiple traces against the same property.

In this spectrum, we build on the widely studied logic LTL_f , or LTL over finite traces (De Giacomo and Vardi 2013), where both the monitored prefixes and their infinitely many suffixes have a finite (yet unbounded) length. Specifically, we focus on the more sophisticated setting of RV-LTL by Bauer, Leucker, and Schallhart (2010), where the verdict of the monitor of a property ψ at each point in time is one of four possible values: *currently satisfied* (ψ is satisfied now but may be violated in a future continuation), *permanently satisfied* (ψ is satisfied now and will necessarily stay so), and the two complementary values of *permanent* and *current violation*. Thus, RV-LTL monitoring is at least as hard as satisfiability and validity. For propositional LTL_f , an RV-LTL monitor can be constructed from the deterministic finite-state automaton (DFA) corresponding to the property, by labeling each DFA state by an RV-LTL value, depending on whether it is final and can reach final and non-final states (Maggi et al. 2011; De Giacomo et al. 2022).

Prior work on LTL_f monitoring like the above has often focused on propositional traces, where states are described by propositional interpretations. This coarse-grained representation creates a large abstraction gap w.r.t. dynamic systems, which often generate *data-aware* traces where states contain richer objects from an infinite domain, e.g., strings or numbers. On the other hand, for richer logics partial methods have been devised, but without guarantees that the monitoring task can be solved, e.g. (Reger, Cruz, and Rydeheard 2015). Aiming to bridge this gap, we (i) introduce an extension of LTL_f , called $ALTL_f$, which combines linear-time operators with linear arithmetic constraints over data variables *within and across states*, (ii) study how to construct automata-based RV-LTL monitors for such formulae and (iii) identify decidable classes of properties, given that $ALTL_f$ monitoring is in general undecidable. The next example illustrates the challenge that we address.

Example 1. Consider the monitoring of suspicious bidding patterns in an on-line auction. The process is supposed

(but not guaranteed) to execute according to the following transition system:



There are three variables: the current price p (i.e., the last bid); the id of the last bidder b , and a timer t . In the edge labels, a primed variable (e.g. p') indicates the next value of the respective variable that is written by the transition (e.g. p). The transition `set` fixes p to some base price and initializes the timer to the duration T of a round. Then, either the timer is decreased by transition `dec`, or a bidder increases the price, upon which the timer is reset to T . Finally, `exp` terminates the auction when the timer expires. (We assume that all variables not explicitly written in transition labels keep their value.) By adding a variable s for the state, the process can be encoded in ALTL_f:

$$(s = \text{init}) \wedge G(\psi_{\text{set}} \vee \psi_{\text{dec}} \vee \psi_{\text{bid}} \vee \psi_{\text{exp}}) \quad (\text{P})$$

where, e.g., $\psi_{\text{set}} = (s = \text{init} \wedge s' = \text{act} \wedge t' = T \wedge p' > 0)$, and other transitions are encoded similarly. A first task could then be to monitor formula P to check whether the actual execution deviates from this specification.

In addition, during an auction, it is desirable to identify users that exhibit shilling behavior, i.e., they drive up the auction price for the seller. Given a user u , consider the following behavior patterns (Xu and Cheng 2007): (OB) Overbidding: u places an overbid (i.e., increasing the price by at least 20%) right before the timer expires. This can be expressed as $OB_u := F(b' = u \wedge t \leq \epsilon \wedge p' \geq p \cdot 1.2)$, for some small ϵ . (AU) Aggressive underbidding: underbidding means that the price increase is less than 3\$, and aggressive that the time gap between two bids is small (for simplicity, we assume here in the next state). This can be expressed as $AU_u := G(b \neq u \wedge b > 0 \rightarrow p' \geq p + 3 \wedge b = u)$. Both (OB) and (AU) likely imply shilling. The following is an example process execution (also called a trace) with $T = 2$.

s	$init$	act	act	act	act	act	act	act
t	0	2	2	2	1	2	2	2
p	0	10	30	32	32	36	40	50
b	0	0	1	2	2	3	1	2
OB_2	CV	CV	CV	CV	CV	CV	CV	PS
AU_2	CS	CS	CS	CS	CS	CS	PV	PV
$shill_2$	CS	CS	CS	CS	CS	CS	CV	PS

The lower part of this table shows monitoring results for the trace, namely the RV values obtained for observing (OB) and (AU) for user 2, as well as for $shill_2 := OB_2 \vee AU_2$. Note that the trace above conforms with the specification P . A further task could be to monitor combinations of such properties, such as $P \wedge OB_2$, which gives the ability of anticipatorily detect violations that cannot be ascribed to single properties, but only to the interplay of multiple properties and the current trace (De Giacomo et al. 2022).

The techniques presented in this paper show how to construct anticipatory monitors for full ALTL_f, to tackle examples like the above. Anticipatory monitoring for plain LTL

has been recognized as highly relevant in business process management (BPM) (Ly et al. 2015; Maggi et al. 2014), but these works ignore the recent trend of modelling data in processes; with the present work we want to mitigate this shortcoming. However, monitoring is also highly relevant in other areas, like program verification (Havelund, Reger, and Rosu 2019) and verification of cyber-physical systems (Geist, Rozier, and Schumann 2014).

Technically, we follow Demri and D’Souza (2007) and consider traces consisting of assignments of a fixed set of variables to numeric values. A starting point is the approach in (Felli, Montali, and Winkler 2022b), where automata-based techniques are used to model check dynamic systems against an extension of LTL_f with linear arithmetic constraints, e.g., “the value of sensor x differs from that of sensor y by 5 units until it exceeds 25”. However, that approach has two prohibitive technical limitations. The first one is operational: their automata are too weak for monitoring as they are faithful only for reachability of final, but not of non-final states. The second shortcoming is semantical: arithmetic conditions can only relate variables within the same state, but not across instants in the trace. For this reason Demri and D’Souza (2007) and, more recently, Geatti, Gianola, and Gigante (2022), allow properties to compare variables across states. Our logic ALTL_f follows this route, enabling variable *lookahead* to reference variable values in different instants within arbitrary linear arithmetic constraints, (cf. e.g. $p' > p$ in Ex. 1). Our approach is thus substantially different from (Demri and D’Souza 2007), which restricts to variable-to-variable/constant comparisons to ensure decidability of satisfiability. Instead, we allow arbitrary linear arithmetic constraints with lookahead, making satisfiability (and monitoring) undecidable even for simple formulae.

Despite the rich research body on monitoring, there are few approaches that are anticipatory, feature full LTL_f, and incorporate arithmetic with lookahead (Falcone et al. 2021). The seminal LOLA approach (D’Angelo et al. 2005) supports arithmetic with lookahead and finite traces, but is not anticipatory. Also Faymonville et al. (2019) and (Basin et al. 2015) have only bounded anticipation. More related is MarQ (Reger, Cruz, and Rydeheard 2015), a strong, anticipatory RV tool for finite traces with relations and arithmetic, but its approach has no decidability guarantees.

We overcome these limitations as follows: (1) we show that ALTL_f monitoring is decidable for linear arithmetic constraints without lookahead, and provide techniques to construct automata-based RV-LTL monitors. (2) we show, for formulae with lookahead, undecidability of ALTL_f satisfiability/monitoring. (3) To mitigate this, we study how the abstract property of *finite summary*, introduced in (Felli, Montali, and Winkler 2022b), can be employed for ALTL_f formulae with lookahead that satisfy this property, proving decidability of monitoring and providing a technique to construct automata-based RV-LTL monitors. (4) We use this general result to show decidability for different concrete classes of formulae enjoying finite summary, obtained by restricting either the language of constraints, or the way these interact with each other via lookahead. (5) A prototype tool witnesses feasibility of our approach.

The structure of the paper reflects Contributions (1)–(5) in this order, after the introduction of ALTL_f and the monitoring problem and a section about the construction of automata for ALTL_f properties afterwards.

Preliminaries

We consider the sorts *int* and *rat* with domains $d(\textit{int}) = \mathbb{Z}$ and $d(\textit{rat}) = \mathbb{Q}$. For a set of variables V , let $V_\sigma \subseteq V$ be those in V of sort σ . We first define arithmetic constraints:

Definition 2. Given a set of variables V , expressions e_σ of sort σ and *constraints* c are defined as follows:

$$\begin{aligned} e_\sigma &:= v_\sigma \mid k_\sigma \mid e_\sigma + e_\sigma \mid k_\sigma \cdot e_\sigma \\ c &:= e_\sigma = e_\sigma \mid e_\sigma \neq e_\sigma \mid e_\sigma < e_\sigma \mid e_\sigma \leq e_\sigma \mid \\ &\quad e_{\textit{int}} \approx_n e_{\textit{int}} \mid e_{\textit{int}} \not\approx_n e_{\textit{int}} \end{aligned}$$

where $k_\sigma \in d(\sigma)$ is a constant, $v_\sigma \in V_\sigma$ a variable of appropriate sort, and \approx_n denotes congruence modulo $n \in \mathbb{N}$.

The set of all constraints over V is denoted by $\mathcal{C}(V)$. E.g., $x < y + z$, and $x - y \neq 2$ are constraints independent of the sort of x, y, z , but $u \approx_3 v + 1$ requires that u and v have sort *int*. We also consider boolean formulas with constraints as atoms; these are in the realm of SMT with linear arithmetic, which is decidable and admits *quantifier elimination* (Presburger 1929): if φ is a formula with free variables $X \cup \{y\}$ and constraint atoms, there is some φ' with free variables X that is logically equivalent to $\exists y. \varphi$, i.e., $\varphi' \equiv \exists y. \varphi$.

From now on, let V be a fixed, finite, non-empty set of variables called *state variables*. An *assignment* α maps every $v \in V$ to a value $\alpha(v)$ in the domain of its sort. The set of *variables with lookahead* \mathcal{V} consists of all $v^{i \dots'}$ such that $v \in V$ is decorated with i prime symbols $'$, for $i \geq 0$, and we say that v has *lookahead* i . Intuitively, $v^{i \dots'}$ with i primes refers to the value of v looking i instants into the future. For instance, v' refers to the value of v in the next step, while v'' is the value of v two steps ahead.

Definition 3. Let \mathcal{L} be the set of properties ψ defined by the following grammar, where $c \in \mathcal{C}(\mathcal{V})$ is a constraint over variables with lookahead:

$$\psi ::= c \mid \psi \wedge \psi \mid \neg \psi \mid X_s \psi \mid \psi \text{ U } \psi$$

Here X_s is the next operator (we use X_w for weak next) and U is the until operator. When needed, we write $\mathcal{L}(V)$ instead of \mathcal{L} to make explicit the set of variables in constraints. The usual transformations apply, namely $\psi_1 \vee \psi_2 \equiv \neg(\psi_1 \wedge \psi_2)$, $X_w \psi \equiv \neg(X_s \neg \psi)$, $F\psi \equiv \top \text{ U } \psi$, $G\psi \equiv \neg F \neg \psi$; moreover $\top \equiv (v = v)$ for some $v \in V$, and $\perp \equiv \neg \top$. A property ψ has *lookahead* m if the maximal lookahead of a variable in ψ is m . E.g., all properties in Ex. 1 are in $\mathcal{L}(V)$ for $V = \{p, t, b\}$. Properties $p' > p$ and AU_2 have lookahead 1, but $t \leq 0$ has lookahead 0. Note that lookahead strictly extends the expressiveness of the language: for instance, the simple property $p' > p$ cannot be expressed by using X_s and X_w alone.

Properties are evaluated over *traces*: a trace τ of length $n \geq 1$ is a finite sequence $\alpha_0, \alpha_1, \dots, \alpha_{n-1}$ of assignments with domain V . We write $\tau(i)$ to denote α_i , for $0 \leq i < n$. A constraint c is *well-defined* at instant i of τ if $0 \leq i < n$ and all variables with lookahead j in c satisfy $i + j < n$. The

upper part of the table in Ex. 1 is an example of a trace, and e.g., $p' \geq p$ is well-defined at all but the last instant.

Definition 4. If an expression e is well-defined at instant i of a trace τ of length n , its *evaluation* $[\tau, i](e)$ at instant i is:

$$\begin{aligned} [\tau, i](k) &= k & [\tau, i](e + e_2) &= [\tau, i](e_1) + [\tau, i](e_2) \\ [\tau, i](v^{i \dots'}) &= \tau(i+j)(v) & [\tau, i](k \cdot e) &= k \cdot [\tau, i](e) \end{aligned}$$

where $v^{i \dots'}$ has lookahead j . Then, τ *satisfies* $\psi \in \mathcal{L}$, denoted $\tau \models \psi$, iff $\tau, 0 \models \psi$ holds, where, for $0 \leq i$:

$$\begin{aligned} \tau, i &\models e_1 \odot e_2 \text{ iff } 0 \leq i < n \text{ and either} \\ &\quad - e_1 \odot e_2 \text{ is not well-defined for } \tau \text{ and } i, \text{ or} \\ &\quad - [\tau, i](e_1) \odot [\tau, i](e_2) \text{ holds} \\ \tau, i &\models \psi_1 \wedge \psi_2 \text{ iff } \tau, i \models \psi_1 \text{ and } \tau, i \models \psi_2 \\ \tau, i &\models \neg \psi \text{ iff } \tau, i \not\models \psi \\ \tau, i &\models X_s \psi \text{ iff } i < n - 1 \text{ and } \tau, i+1 \models \psi \\ \tau, i &\models \psi_1 \text{ U } \psi_2 \text{ iff } i < n \text{ and either } \tau, i \models \psi_2, \text{ or} \\ &\quad i < n - 1, \tau, i \models \psi_1 \text{ and } \tau, i+1 \models \psi_1 \text{ U } \psi_2 \end{aligned}$$

For instance, the trace in Ex. 1 satisfies OB_2 but its proper prefixes do not. Note that lookahead variables have a *weak* semantics, as a constraint c holds if it contains a variable $v^{i \dots'}$ that refers to an instant beyond the end of the trace. If instead a strict semantics is desired, one can replace c by $c \wedge X_s \dots X_s \top$, where the number of X_s operators equals the lookahead of c .

We now define our main task, that is to monitor how the satisfaction of a given property changes along a trace, that is, by considering the trace (fragment) τ seen *so far*. As customary (Bauer, Leucker, and Schallhart 2010), we consider the set $RV = \{\text{PS}, \text{CS}, \text{CV}, \text{PV}\}$ of four distinct *monitoring states*: current satisfaction (CS), permanent satisfaction (PS), current violation (CV) and permanent violation (PV).

Definition 5. A property $\psi \in \mathcal{L}$ is in monitoring state $s \in RV$ after a trace τ , written $\tau \models \llbracket \psi = s \rrbracket$, if

- $s = \text{CS}$, $\tau \models \psi$, and $\tau \tau' \not\models \psi$ for some trace τ' ;
- $s = \text{PS}$, $\tau \models \psi$, and $\tau \tau' \models \psi$ for every trace τ' ;
- $s = \text{CV}$, $\tau \not\models \psi$, and $\tau \tau' \models \psi$ for some trace τ' ;
- $s = \text{PV}$, $\tau \not\models \psi$, and $\tau \tau' \not\models \psi$ for every trace τ' .

E.g., $\tau \models \llbracket \psi = \text{CS} \rrbracket$ means that ψ is currently true after τ but there exists a possible continuation of τ (i.e., a trace $\tau \tau'$) after which ψ is false. After a trace, a property ψ is in exactly one possible monitoring state. For instance, the monitoring states in the table in Ex. 1 reflect Def. 5 for the given trace τ , e.g., $\tau \models \llbracket OB_2 = \text{PS} \rrbracket$.

Given τ and ψ , the *monitoring problem* is to compute the state $s \in RV$ s.t. $\tau \models \llbracket \psi = s \rrbracket$. It is *solvable* if one can construct a procedure for ψ that computes the monitoring state for any given trace. Unfortunately, we have that in general:

Theorem 6. *The monitoring problem is not solvable.*

To see this, call a property $\psi \in \mathcal{L}$ *satisfiable* if $\tau \models \psi$ for some (non-empty) trace τ . It can be shown that satisfiability of \mathcal{L} properties with lookahead 1 is undecidable, by a reduction from reachability in 2-counter machines. Now, for a formula of the form $\psi = X_s \psi'$ and a trace τ_0 of length 1, ψ' is satisfiable iff $\tau_0 \models \llbracket \psi = \text{CV} \rrbracket$, so satisfiability of ψ' reduces to monitoring τ_0 against ψ .

In this paper we study when monitoring is solvable. First, we show that without loss of generality one can restrict to

properties with lookahead 1: Indeed, a property with lookahead $m > 1$ can be transformed into one with $m = 1$ by adding fresh variables and extending to these the assignments in the trace in an appropriate way. In general, $m - 1$ fresh book-keeping variables are needed for each variable with lookahead $m > 1$. Intuitively, constraints are then rewritten to mention only consecutive instants, using chains of next operators. E.g., $\psi = G(x'' > x)$ is equivalent to $\widehat{\psi} = G(u' > x \wedge X_w(u = x'))$. The trace $\tau: \{x = 2\}, \{x = 0\}, \{x = 3\}$ satisfies ψ and $\widehat{\tau}: \{x = 2, u = 0\}, \{x = 0, u = 3\}, \{x = 3, u = 0\}$ satisfies $\widehat{\psi}$, with $\widehat{\tau}$ obtained from τ by assigning to u the subsequent value of x , if it exists. A detailed proof of the next result can be found in (Felli et al. 2022).

Lemma 7. Let $\psi \in \mathcal{L}(V)$ have lookahead m . There is some $\psi_1 \in \mathcal{L}(V \cup X)$ with lookahead 1 for a set of fresh variables X such that for every trace τ over V there is a trace $\widehat{\tau}$ over $V \cup X$ satisfying $\tau \models \psi$ iff $\widehat{\tau} \models \psi_1$, where $|\widehat{\tau}| = |\tau|$ and $\widehat{\tau}(i)$ agrees with $\tau(i)$ on V , for all i .

Automata for properties

In the sequel, we will develop first monitoring techniques for \mathcal{L} properties without lookahead, and then with lookahead 1. Although the former case is simpler and could be addressed by an ad-hoc approach resembling techniques for propositional LTL_f (Maggi et al. 2011), we adopt a uniform approach for both cases. To this end, we first present the construction of an automaton that represents a given property ψ ; this will be crucial in the later sections. As a preprocessing step, two transformations are applied to the property:

(1) We assume ψ in negation normal form; so it may also contain \vee , G , and X_w , so far considered syntactic sugar. Note that for constraints with lookahead, negation cannot be pushed to the constraint level, e.g., $\neg(x' = x)$ is not equivalent to $x' \neq x$ as the latter is satisfied by any trace of length 1 but the former is not. In the construction, we thus treat constraints with and without negation separately. For clarity, we write $neg(c)$ to flip the comparison operator in c : $neg(t_1 = t_2) := t_1 \neq t_2$, $neg(t_1 < t_2) := t_2 \leq t_1$, $neg(t_1 \approx_n t_2) := t_2 \not\approx_n t_1$, and vice versa, for all t_1, t_2 .

(2) We want a monitor to return the monitoring state for the trace (prefix) τ seen so far, without depending on future variable assignments. However, the evaluation of a constraint with lookahead depends on the next, not yet seen, assignment (see Def. 4). To avoid this counterintuitive peculiarity and make the monitoring task clearer, we transform constraints with lookahead 1 into constraints with lookback 1.

To this end, we consider variables v_{pre} and v_{cur} , for all $v \in V$. For a constraint c , let $back(c)$ be defined as (a) if c has no lookahead, $back(c)$ is obtained from c by replacing v with v_{cur} , and (b) if c has lookahead, $back(c) := X_w \bar{c}$, with \bar{c} obtained from c by replacing v with v_{pre} and v' with v_{cur} ; for all $v \in V$. We denote the property where each constraint c is replaced by $back(c)$ by ψ_{back} , and we write $V_{pre} = \{v_{pre} \mid v \in V\}$ and the same for V_{cur} . Then ψ and ψ_{back} are equivalent, in the sense that $1, \tau \models \psi$ iff $1, \tau \models \psi_{back}$, where the evaluation in Def. 4 is extended to variables V_{pre} and V_{cur} as follows: $[\tau, i](v_{cur}) = \tau(i)(v)$ $0 \leq i \leq n$ and $[\tau, i](v_{pre}) = \tau(i-1)(v)$ for $0 < i \leq n$. The

proof is straightforward and it is omitted.

Example 8. Property $\psi = G(x' > x) \wedge F(x = 2)$ is equivalent to $\psi_{back} = G(X_w(x_{cur} > x_{pre})) \wedge F(x_{cur} = 2)$, it demands that x is increasing, and at some point has value 2.

Given ψ_{back} , we build an NFA $\mathcal{N}_{\psi_{back}}$ using an auxiliary function δ , as in (De Giacomo and Vardi 2013). Let C be the set of constraints in ψ_{back} , and $C^\pm = C \cup \{neg(c) \mid c \in C\}$; and λ be an auxiliary proposition used to mark the last element of a trace. The alphabet of the NFA will be $\Sigma = 2^{C^\pm}$, i.e., a symbol is a set of constraints. For the construction we also use its extension $\Sigma_\lambda = 2^{C^\pm \cup \{\lambda, \neg\lambda\}}$. Let $\varsigma \in \Sigma_\lambda$ be *satisfiable* if $\{\lambda, \neg\lambda\} \not\subseteq \varsigma$, and the conjunction of constraints in ς is satisfiable. The input of δ is a (quoted) property $\phi \in \mathcal{L} \cup \{\top, \perp\}$ over $V_{pre} \cup V_{cur}$. The output of δ is a set of pairs (ϕ'', ς) where $\phi'' \in \mathcal{L} \cup \{\top, \perp\}$ is again a (quoted) property over $V_{pre} \cup V_{cur}$ and $\varsigma \in \Sigma_\lambda$. For two sets of such pairs R_1, R_2 , let $R_1 \otimes R_2 = \{(\psi_1 \wedge \psi_2, \varsigma_1 \cup \varsigma_2) \mid (\psi_1, \varsigma_1) \in R_1, (\psi_2, \varsigma_2) \in R_2 \text{ and } \varsigma_1 \cup \varsigma_2 \text{ is satisfiable}\}$, where we simplify $\psi_1 \wedge \psi_2$ if possible; and let $R_1 \oplus R_2$ be defined in the same way, replacing con- with disjunction.

Intuitively, $(\phi'', \varsigma) \in \delta(\phi)$ expresses that when ϕ is the property to evaluate, if ς is the current symbol of the word $w \in \Sigma^+$ that is read, then ϕ'' is the (sub)property yet to satisfy to determine whether w satisfies ϕ .

Definition 9. For $\psi \in \mathcal{L} \cup \{\top, \perp\}$, δ is as follows:

$$\begin{aligned} \delta(\top) &= \{(\top, \emptyset)\} \text{ and } \delta(\perp) = \{(\perp, \emptyset)\} \\ \delta(c) &= \{(\top, \{c\}), (\perp, \{neg(c)\})\} \\ \delta(\neg c) &= \{(\perp, \{c\}), (\top, \{neg(c)\})\} \\ \delta(\psi_1 \wedge \psi_2) &= \delta(\psi_1) \otimes \delta(\psi_2) \\ \delta(\psi_1 \vee \psi_2) &= \delta(\psi_1) \oplus \delta(\psi_2) \\ \delta(X_s \psi) &= \{(\psi, \{\neg\lambda\}), (\perp, \{\lambda\})\} \\ \delta(X_w \psi) &= \{(\psi, \{\neg\lambda\}), (\top, \{\lambda\})\} \\ \delta(G\psi) &= \delta(\psi) \otimes \delta(X_w G\psi) \\ \delta(\psi_1 \text{ U } \psi_2) &= \delta(\psi_2) \oplus (\delta(\psi_1) \otimes \delta(X_s(\psi_1 \text{ U } \psi_2))). \end{aligned}$$


While the symbol λ is needed for defining δ , we can omit it from the NFA, and define \mathcal{N}_ψ as follows:

Definition 10. Given a property $\psi \in \mathcal{L}$, we define the NFA as $\mathcal{N}_\psi = (Q, \Sigma, \rho, q_0, Q_F)$ where $q_0 = \psi_{back}$ is the initial state, $Q_F = \{\top, q_+\}$ is the subset of final states, q_- is an additional non-final state, and Q, ρ are the smallest sets such that $q_0, q_F, q_+, q_- \in Q$ and whenever $q \in Q \setminus \{q_+, q_-\}$ and $(q', \varsigma) \in \delta(q)$ then $q' \in Q$ and

- (i) if $\lambda \notin \varsigma$ then $(q, \varsigma \setminus \{\neg\lambda\}, q') \in \rho$, and
- (ii) whenever $\lambda \in \varsigma$, if $q' = \top$ then $(q, \varsigma \setminus \{\lambda\}, q_+) \in \rho$, and if $q' = \perp$ then $(q, \varsigma \setminus \{\lambda\}, q_-) \in \rho$.

The following example illustrates the construction.

Example 11. For $\psi_{back} = GX_w(x_{cur} \geq x_{pre}) \wedge F(x_{cur} = 2)$ from Ex. 8, we get the following NFA $\mathcal{N}_{\psi_{back}}$:



A word $w = \varsigma_0, \varsigma_1, \dots, \varsigma_{n-1} \in \Sigma^+$ is *well-formed* if there are no (negated) constraints in ς_0 that mention V_{pre} . In Ex. 11 one can see that all words accepted by the NFA are well-formed (this is guaranteed by the translation to ψ_{back}). We next define the notion of consistency to relate words $w \in \Sigma^+$ as above (where each ς_i is a set of constraints) with those traces $\tau = \alpha_0, \alpha_1, \dots, \alpha_{n-1}$ s.t., at each step, the assignment α_i satisfies ς_i , i.e., the traces fitting w . To that end, for assignments α and α' with domain V , define for all $v \in V$ the assignment $[\alpha \otimes \alpha']$ with domain $V_{cur} \cup V_{pre}$ as $[\alpha \otimes \alpha'](v_{pre}) = \alpha(v)$ and $[\alpha \otimes \alpha'](v_{cur}) = \alpha'(v)$.

Definition 12. A well-formed word $\varsigma_0, \varsigma_1, \dots, \varsigma_{n-1} \in \Sigma^+$ is *consistent* with a trace $\alpha_0, \alpha_1, \dots, \alpha_{n-1}$ if $\alpha_0 \models \varsigma_0$ and $[\alpha_{i-1} \otimes \alpha_i] \models \varsigma_i$ for all $0 < i < n$.

For instance, the word $\{x_{cur} \neq 2\} \{x_{cur}=2, x_{cur} \geq x_{pre}\}$ is consistent with trace $\{x=1\} \{x=2\}$, but not with $\{x=3\} \{x=2\}$. The key property of \mathcal{N}_ψ is stated in the following theorem: a word w is accepted iff all the traces captured by w satisfy ψ .

Theorem 13. Given $\psi \in \mathcal{L}$, a trace τ and a well-formed word $w \in \Sigma^+$ consistent with τ , $\mathcal{N}_{\psi_{back}}$ accepts w iff $\tau \models \psi$.


However, $\mathcal{N}_{\psi_{back}}$ is not deterministic, which is inconvenient for monitoring. Thus, we next build an equivalent DFA $\mathcal{D}_{\psi_{back}}$ (exponentially larger in the worst case), using a subset construction, and a restricted alphabet: Let $C_{cur} \subseteq C$ be the subset of constraints in ψ_{back} that mention only V_{cur} (but not V_{pre}). Let the alphabet $\Theta \subseteq \Sigma$ consist of all maximal satisfiable subsets of C^\pm , and $\Theta_{cur} \subseteq \Sigma$ consist of all maximal satisfiable subsets of C_{cur}^\pm .

Definition 14. For $\mathcal{N}_{\psi_{back}} = (Q, \Sigma, \varrho, q_0, Q_F)$, let $\mathcal{D}_{\psi_{back}} = (Q, \Theta \cup \Theta_{cur}, \Delta, \{q_0\}, Q_F)$ where $Q = 2^Q$, Q_F consists of all $P \subseteq Q$ such that $P \cap Q_F \neq \emptyset$, and the transition function is given by $\Delta(\{q_0\}, \varsigma_0) = \{q' \mid (q_0, \varsigma', q') \in \varrho \text{ and } \varsigma' \subseteq \varsigma_0\}$ for all $\varsigma_0 \in \Theta_{cur}$, and $\Delta(P, \varsigma) = \{q' \mid q \in P, (q, \varsigma', q') \in \varrho \text{ and } \varsigma' \subseteq \varsigma\}$ for all $\varsigma \in \Theta$ and $P \neq \{q_0\}$.

Note that Δ is defined either (i) for the initial state $\{q_0\}$ and $\varsigma \in \Theta_{cur}$, i.e., a constraint set not mentioning V_{pre} , or (ii) for a non-initial state and $\varsigma \in \Theta$. This distinction ensures that $\mathcal{D}_{\psi_{back}}$ accepts only well-formed words, and \mathcal{D}_ψ is *deterministic* in that for every $w \in \Theta_{cur} \Theta^*$ there is a unique state P of $\mathcal{D}_{\psi_{back}}$ such that $\{q_0\} \rightarrow_w^* P$. \mathcal{D}_ψ is equivalent to \mathcal{N}_ψ in the following sense:

Lemma 15. Given a trace τ , \mathcal{D}_ψ accepts a word consistent with τ iff \mathcal{N}_ψ accepts a word consistent with τ .

Example 16. We build the DFA for the property ψ_{back} and the respective NFA from Ex. 11. Below, Θ_{cur} consists of $\{x_{cur} = 2\}$ and $\{x_{cur} \neq 2\}$, and Θ of all four combinations of $x_{cur} = 2$ or $x_{cur} \neq 2$ with $x_{cur} \geq x_{pre}$ or $x_{cur} < x_{pre}$:




Here A corresponds to $\{q_0\}$, B to $\{q_2, q_+\}$ (equivalent to $\{q_2, q_+, \perp\}$), C to $\{q_1, \perp\}$, and D to $\{\perp\}$.


Monitoring properties without lookahead

In this section we show that monitoring properties without lookahead is solvable, and a monitoring structure is given by the DFA $\mathcal{D}_{\psi_{back}} = (\mathcal{Q}, \Theta, \delta, \{q_0\}, \mathcal{Q}_F)$ (note that without lookahead, Θ_{cur} and Θ as defined in the last section coincide). We illustrate the construction on an example.

Example 17. Let $\psi = (y \geq 0) \cup (x > y \wedge G(x > y))$. Then ψ_{back} is as ψ but where x is replaced by x_{cur} and y by y_{cur} . As there is no lookahead, no confusion can arise, so we write x for x_{cur} and y for y_{cur} . We get the following NFA $\mathcal{N}_{\psi_{back}}$ with alphabet $\Sigma = 2^{C^\pm}$ for $C = \{x > y, y \geq 0\}$:



The next DFA $\mathcal{D}_{\psi_{back}}$ is equivalent; its alphabet consists of all maximal satisfiable subsets of C^\pm , i.e., $\{y \geq 0, x > y\}$, $\{y \geq 0, x \leq y\}$, $\{y < 0, x > y\}$, and $\{y < 0, x \leq y\}$.



Here $A = \{q_1\}$, $B = \{q_2\}$, $C = \{q_1, q_2\}$, and $D = \emptyset$.

Next, we formalize how to use $\mathcal{D}_{\psi_{back}}$ as a monitor. We call a state P' *reachable* from P if $P \rightarrow_w^* P'$ for some w .

Theorem 18. If a word $w \in \Theta^+$ is consistent with a trace τ and P is the state in $\mathcal{D}_{\psi_{back}}$ such that $\{q_0\} \rightarrow_w^* P$,

- $\tau \models \llbracket \psi = \text{PS} \rrbracket$ if only final states are reachable from P ,
- $\tau \models \llbracket \psi = \text{CS} \rrbracket$ if $P \in \mathcal{Q}_F$ but $P \rightarrow_w^* P'$ for some $P' \notin \mathcal{Q}_F$,
- $\tau \models \llbracket \psi = \text{CV} \rrbracket$ if $P \notin \mathcal{Q}_F$ but $P \rightarrow_w^* P'$ for some $P' \in \mathcal{Q}_F$,
- $\tau \models \llbracket \psi = \text{PV} \rrbracket$ if no final state is reachable from P .

Proof. By Lems. 15 and 13, $\tau \models \psi$ iff P is final. (Case PS:) P has no path to a non-final state. Towards contradiction, assume τ' exists s.t. $\tau\tau' \not\models \psi$. By Lems. 13 and 15, $\mathcal{D}_{\psi_{back}}$ does not accept a word consistent with $\tau\tau'$. First, we show that for every non-empty trace τ there is a unique word $w \in \Theta_{cur} \Theta^*$ consistent with τ . Let $\tau = \langle \alpha_0, \dots, \alpha_{n-1} \rangle$, and C be the constraints in ψ_{back} . Since for every constraint $c \in C_{cur}$, either $\alpha \models c$ or $\alpha \models \text{neg}(c)$, and Θ_{cur} contains all maximal satisfiable subsets of $C_{cur} \cup C_{cur}^-$, there must be exactly one subset $\varsigma_0 \in \Theta_{cur}$ s.t. $\alpha \models \bigwedge \varsigma_0$. Similarly, $[\alpha_i \otimes \alpha_{i+1}]$ must model either c or $\text{neg}(c)$ for all constraints in c , so there must be exactly one $\varsigma_i \in \Theta$ s.t. $[\alpha_i \otimes \alpha_{i+1}] \models \bigwedge \varsigma_i$, for all $0 < i \leq n-1$. Thus, the word $w = \langle \varsigma_0, \dots, \varsigma_{n-1} \rangle$ is unique and consistent with τ .

With this result at hand, by determinism, we can write $\bar{w} = ww'$ s.t. $\{q_0\} \rightarrow_w^* P \rightarrow_{w'}^* P'$, contradicting that only final states are reached from P . (Case CS:) let $P \rightarrow_w^* P'$ for a non-final state P' . Sets in Θ are satisfiable, so there is a trace τ' s.t. w' is consistent with τ' , and ww' with $\tau\tau'$. By the result above, ww' is the unique word in Θ^* consistent with $\tau\tau'$, and not accepted. By Lems. 15 and 13, $\tau\tau' \not\models \psi$, so $\tau \models \llbracket \psi = \text{CS} \rrbracket$. The other cases are dual. \square

Thm. 18 implies that monitoring is solvable, and the monitoring state only depends on the DFA state matching a trace.

Example 19. Consider the property of Ex. 17. The following example trace is interleaved with the state reached by the corresponding word in $\mathcal{D}_{\psi_{back}}$, and the monitoring state:

$$1 \left\{ \begin{array}{l} x=0 \\ y=0 \end{array} \right\} \text{CV} \quad 1 \left\{ \begin{array}{l} x=0 \\ y=3 \end{array} \right\} \text{CV} \quad 1 \left\{ \begin{array}{l} x=4 \\ y=3 \end{array} \right\} \text{CS} \quad 12 \left\{ \begin{array}{l} x=0 \\ y=3 \end{array} \right\} \text{CV} \quad 1 \left\{ \begin{array}{l} x=0 \\ y=-1 \end{array} \right\} \text{CS} \quad 2$$

Each DFA state corresponds to a monitoring state, as indicated by the colors in Ex. 17: non-final states are CV if a final state is reachable (yellow), PV otherwise (red); final states are CS if a non-final state is reachable (blue), PS otherwise.

Solvable cases for properties with lookahead

In this section we study the more involved case of properties with lookahead 1. The next example shows that the approach from the last section does not work for this case:

Example 20. Consider the property and DFA from Exa. 16 and the following trace, interleaved with the reached state in $\mathcal{D}_{\psi_{back}}$ as well as the monitoring state for every prefix:

$$A \left\{ \begin{array}{l} x=0 \end{array} \right\} \text{CV} \quad C \left\{ \begin{array}{l} x=1 \end{array} \right\} \text{CV} \quad C \left\{ \begin{array}{l} x=3 \end{array} \right\} \text{PV} \quad C \left\{ \begin{array}{l} x=4 \end{array} \right\} \text{PV}$$

The reached DFA state is the same for every prefix, even though the monitoring state changes.

Note that while the state P reached by a trace τ in $\mathcal{D}_{\psi_{back}}$ can still tell whether τ satisfies ψ , reachability of (non)final states from P is no longer sufficient to predict whether the trace can be extended to satisfy ψ . We therefore develop an enhanced monitoring structure. To simplify notation, we assume that ψ is the result of the $back(\cdot)$ transformation, i.e., a property over variables $V_{pre} \cup V_{cur}$. Moreover, we assume throughout this section that \mathcal{D}_{ψ} is a DFA for ψ of the form $\mathcal{D}_{\psi} = (\mathcal{Q}, \Theta \cup \Theta_{cur}, \Delta, \{q_0\}, \mathcal{Q}_F)$.

The idea is to abstract the variable configurations that are relevant for verification by a set of formulas. This abstraction requires some additional notation: First, let $V_0 = \{v_0 \mid v \in V\}$ be a copy of V , it will act as placeholders for the initial values of the variables V . We write φ_{init} for the formula $\varphi_{init} = \bigwedge_{v \in V} v = v_0$. For a formula φ with free variables $V_0 \cup V$, let $\varphi(\bar{U})$ denote the formula where \bar{V} is replaced by \bar{U} , but V_0 is not replaced (with denote by \bar{V} a fixed ordering of V as a vector). For a set of constraints C with free variables $V_{pre} \cup V_{cur}$, let $C(\bar{U}, \bar{V})$ denote the constraints where \bar{V}_{pre} is replaced by \bar{U} and \bar{V}_{cur} is replaced by \bar{V} .

Definition 21. For a formula φ with free variables $V_0 \cup V$ and a set of constraints C with free variables $V_{pre} \cup V_{cur}$, let $update(\varphi, C) = \exists \bar{U}. \varphi(\bar{U}) \wedge \bigwedge C(\bar{U}, \bar{V})$.

Both φ and $update(\varphi, C)$ have free variables $V_0 \cup V$. Although $update$ introduces existential quantifiers, there is an equivalent quantifier-free formula for $update(\varphi, C)$: linear arithmetic has quantifier elimination. E.g., for $\varphi = (x = x_0 \wedge x_0 \neq 2)$ and $C = \{x_{cur} \geq x_{pre}, x_{cur} = 2\}$, we have $update(\varphi, C) = \exists u. u = x_0 \wedge x_0 \neq 2 \wedge x \geq u \wedge x = 2$, which is equivalent to $x_0 < 2 \wedge x = 2$.

In order to design a monitoring structure, we build a *constraint graph* which, intuitively, connects in a graph all variable dependencies (described by formulas) emerging from a state P_0 in the DFA, and where V_0 acts as placeholder for the initial variable values. The constraint graph is parameterized




Figure 1: Constraint graph $CG_{\psi}(C)$.

by an equivalence relation \sim on formulas used to reduce its size. Our correctness proof will require further properties of \sim but a default choice is the logical equivalence relation \equiv .

Definition 22. A *constraint graph* $CG_{\psi}(P_0, \sim)$ for \mathcal{D}_{ψ} and a DFA state $P_0 \in \mathcal{Q}$ is a triple $\langle S, s_0, \gamma, S_F \rangle$ where the node set S consists of tuples (P, φ) of a state $P \in \mathcal{Q}$ and a formula φ with free variables $V_0 \cup V$, and $\gamma \subseteq S \times \Sigma \times S$. Then, S and γ are inductively defined as follows:

- (i) $s_0 = (P_0, \varphi_{init})$ is the initial node and $s_0 \in S$;
- (ii) if $s = (P, \varphi) \in S$ and $P \xrightarrow{\sigma} P'$ in \mathcal{D}_{ψ} such that $update(\varphi, \sigma)$ is satisfiable, there is some $s' = (P', \varphi') \in S$ with $\varphi' \sim update(\varphi, \sigma)$, and $s \xrightarrow{\sigma} s'$ is in γ ;
- (iii) $(P, \varphi) \in S$ is in the set of final nodes S_F iff $P \in \mathcal{Q}_F$.

We simply write $CG_{\psi}(P_0)$ for $CG_{\psi}(P_0, \equiv)$. The constraint graph $CG_{\psi}(C)$ for the DFA $\mathcal{D}_{\psi_{back}}$ from Ex. 16 is shown in Fig. 1. For readability, edge labels are combined; e.g. the two edges from B to D labeled $\{x_{cur} < x_{pre}, x_{cur} = 2\}$ and $\{x_{cur} < x_{pre}, x_{cur} \neq 2\}$ are combined to one edge labeled $\{x_{cur} < x_{pre}\}$. The construction in Def. 22 need not terminate as infinitely many formulas may occur, but we will show in the next section that termination is guaranteed for several relevant cases.

Definition 23. For ψ with DFA \mathcal{D}_{ψ} , a state P_0 in \mathcal{D}_{ψ} , and $G := CG_{\psi}(P_0, \sim)$ with node set S , let

$$\text{FSAT}(G) = \exists \bar{V}. \bigvee \{ \varphi \mid (P, \varphi) \in S \text{ is a final node} \}$$

$$\text{FUNS}(G) = \exists \bar{V}. \bigvee \{ \varphi \mid (P, \varphi) \in S \text{ is a non-final node} \}$$

Intuitively, $\text{FSAT}(G)$ expresses a condition on the variable values when a final state in \mathcal{D}_{ψ} is reachable from P_0 ; dually, FUNS gives a condition for reachability of non-final states. This leads to the following monitoring procedure:

- 1: **procedure** MONITOR(ψ, τ)
- 2: compute ψ_{back} and $\mathcal{D}_{\psi_{back}}$ ▷ or take from cache
- 3: $w \leftarrow$ word in $\Theta_{cur} \Theta^*$ consistent with τ
- 4: $P \leftarrow$ state in $\mathcal{D}_{\psi_{back}}$ such that $\{q_0\} \rightarrow_w^* P$
- 5: $G \leftarrow CG_{\psi}(P, \sim)$ ▷ or take from cache
- 6: $\alpha_n \leftarrow$ last assignment in τ
- 7: **if** P accepting in $\mathcal{D}_{\psi_{back}}$ **then**
- 8: **return** (CS **if** $\alpha_n \models \text{FUNS}(G)$ **else** PS)
- 9: **else return** (CV **if** $\alpha_n \models \text{FSAT}(G)$ **else** PV)

Note that the monitoring structures (DFA, constraint graphs for DFA states, FSAT, FUNS) can be computed once and for all upfront. This is particularly useful when monitoring multiple traces against the same property ψ . In this case, the procedure above should be implemented in a *lazy* way, where the structures are computed when needed and cached for later use. At monitoring time, given trace τ , one thus only

needs to find the DFA state corresponding to τ , and evaluate the pre-computed formula with the last assignment of τ .

Example 24. For $\psi_{back} = \text{GX}_w(x_{cur} \geq x_{pre}) \wedge F(x_{cur} = 2)$ from Exa. 16, by the graph in Fig. 1 we have $\text{FSAT}(C) = \exists x.(x_0 \leq 2 \wedge x \geq 2) \vee (x_0 \leq 2 \wedge x = 2)$, equivalent to $x_0 \leq 2$. We show how the procedure distinguishes two prefixes of the trace from Exa. 20 with different monitoring state:

- Trace $\tau_1 = \langle \{x=0\} \rangle$ corresponds to $w = \langle \{x_{cur} \neq 2\} \rangle$, leading to the non-accepting state C . As $\{x=0\} \models \text{FSAT}(G_C)(V) = x \leq 2$, we have $\tau_1 \models \llbracket \psi = \text{CV} \rrbracket$.
- Trace $\tau_3 = \langle \{x=0\}, \{x=1\}, \{x=3\} \rangle$ matches $\langle \{x_{cur} \neq 2\}, \{x_{cur} \neq 2, x_{cur} \geq x_{pre}\}^2 \rangle$, leading again to state C . However, $\{x=3\} \not\models x \leq 2$, so $\tau_3 \models \llbracket \psi = \text{PV} \rrbracket$.

Correctness. To prove correctness, we need some additional notions. First, we define *history constraints* to capture the formulas obtained from stacked *update* operations:

Definition 25. For $w = s_0, \dots, s_{n-1}$ in Θ^* , the *history constraint* $h(w)$ is given by $h(w) = \varphi_{init}$ if $n = 0$, and if $n > 0$ then $h(w) = \text{update}(h(\langle s_0, \dots, s_{n-2} \rangle), s_{n-1})$.

There is the correspondence between satisfying assignments of $h(w)$ and traces that satisfy the constraints in w :

Lemma 26. For $w \in \Theta^*$ of length n , $h(w)$ is satisfied by assignment ν iff $\langle \emptyset \rangle \cdot w$ is consistent with a trace $\alpha_0, \dots, \alpha_n$ such that $\nu(v_0) = \alpha_0(v)$ and $\nu(v) = \alpha_n(v)$ for all $v \in V$.

Next, we define a *history set* to capture all variable configurations that are relevant for monitoring. More precisely, it collects pairs of history constraints for words and the state that they lead to, but where equivalent pairs are eliminated:

Definition 27. A *history set* for \mathcal{D}_ψ is a minimal set Φ of pairs (P, φ) of a state $P \in \mathcal{Q}$ and a formula φ such that for all P' and $w \in \Theta_{cur} \Theta^*$ with $P \rightarrow_w^* P'$ in \mathcal{D}_ψ , there is a $(P', \varphi) \in \Phi$ s.t. $h(w) \equiv \varphi$.

If Φ is finite, it represents an abstract representation of the relevant variable configurations and can be used to build the states of the constraint graph. For succinctness, we pair a history set with an equivalence relation.

Definition 28. A *summary* for \mathcal{D}_ψ is a pair (Φ, \sim) of a history set Φ for \mathcal{D}_ψ and equivalence relation \sim on Φ such that

- (1) \sim contains \equiv on formulas in Φ and is decidable,
- (2) for all $(q, \varphi), (q, \psi) \in \Phi$ such that $\varphi \sim \psi$, (a) if $\alpha \models \varphi$, there is some α' such that $\alpha' \models \psi$ and $\alpha(u) = \alpha'(u)$ for all $u \in V_0$, and (b) for all transitions $q \xrightarrow{s} q'$, $\text{update}(\varphi, s) \sim \text{update}(\psi, s)$.

The summary (Φ, \sim) is *finite* if \sim has finitely many equivalence classes, and ψ has *finite summary* if a finite summary exists for \mathcal{D}_ψ . For a word $w = s_1, \dots, s_n$, we write $(P_0, \varphi_{init}) \rightarrow_w^* (P, \varphi)$ if there is a path $(P_0, \varphi_{init}) \xrightarrow{s_1} \dots \xrightarrow{s_n} (P, \varphi)$ in $\text{CG}_\psi(P_0, \sim)$. The next lemma states that a path in the CG emulates a path in \mathcal{D}_ψ and leads to a state that is equivalent to a history constraint.

Lemma 29. Let (Φ, \sim) be a summary for \mathcal{D}_ψ . (a) If $\text{CG}_\psi(P_0, \sim)$ has a path $(P_0, \varphi_{init}) \rightarrow_w^* (P, \varphi)$ then $P_0 \rightarrow_w^* P$ in \mathcal{D}_ψ and $\varphi \sim h(w)$ is satisfiable, and (b) if $P_0 \rightarrow_w^* P$ and $h(w)$ is satisfiable then there is a path $(P_0, \varphi_{init}) \rightarrow_w^* (P, \varphi)$ for some φ such that $\varphi \sim h(w)$.

At this point we are ready to prove our main theorem:

Theorem 30. Let (Φ, \sim) be a summary for \mathcal{D}_ψ . Given a DFA \mathcal{D}_ψ for ψ and a trace τ , let $w \in \Theta_{cur} \Theta^*$ be the word consistent with τ and P the \mathcal{D}_ψ state s.t. $\{q_0\} \rightarrow_w^* P$. For $G := \text{CG}_\psi(P, \sim)$ the constraint graph from P ,

- if $P \in P_F$ then $\tau \models \llbracket \psi = \text{CS} \rrbracket$ if $\alpha_n \models \text{FUNS}(G)(\bar{V})$, and $\tau \models \llbracket \psi = \text{PS} \rrbracket$ otherwise,
- if $P \notin P_F$ then $\tau \models \llbracket \psi = \text{CV} \rrbracket$ if $\alpha_n \models \text{FSAT}(G)(\bar{V})$, and $\tau \models \llbracket \psi = \text{PV} \rrbracket$ otherwise.

Proof. Note that the word w that is consistent with τ is unique (cf. proof of Thm. 18). Consider the case where $P \in P_F$, so that $\tau \models \psi$ by Lem. 13 and 15. First, suppose $\alpha_n \models \text{FUNS}(G)(\bar{V})$, so $\alpha_n \models (\exists \bar{V}. \varphi)(\bar{V})$ for some non-final node (P', φ) of G . Let u be a word s.t. there is a path $(P, \varphi_{init}) \rightarrow_u^* (P', \varphi)$ in $\text{CG}_\psi(P)$. By Lem. 29, $P \rightarrow_u^* P'$ in \mathcal{D}_ψ and $\varphi \sim h(u)$. By Def. 28, $\alpha_n \models (\exists \bar{V}. h(u))(\bar{V})$, so there is an assignment ν with domain $V_0 \cup V$ s.t. $\nu(v_0) = \alpha_n(v)$ for all $v \in V$ and $\nu \models h(u)$. By Lem. 26, $\langle \emptyset \rangle \cdot u$ is consistent with some trace $\tau' = \langle \alpha'_0, \dots, \alpha'_k \rangle$ s.t. $\nu(v_0) = \alpha'_0(v)$ for all $v \in V$. Thus, α_n and α'_0 coincide, and wu is consistent with the trace $\tau'' = \langle \alpha_0, \dots, \alpha_n, \alpha'_1, \dots, \alpha'_k \rangle$ as w is consistent with τ and u with τ' . As $\{q_0\} \rightarrow_{wu}^* P'$ is not accepting, by Lem. 13, $\tau'' \not\models \psi$, hence $\tau \models \llbracket \psi = \text{CS} \rrbracket$.

Second, let $\alpha_n \not\models \text{FUNS}(G)(\bar{V})$. Towards a contradiction, suppose there is a trace $\tau' = \langle \alpha'_1, \dots, \alpha'_k \rangle$ s.t. $\tau\tau' \not\models \psi$. Again, there is a unique word $u \in \Theta^*$ s.t. wu is consistent with $\tau\tau'$; let P' the state s.t. $\{q_0\} \rightarrow_{wu}^* P'$. By Lem. 13 and 15, the state P' is not final. As wu is consistent with $\tau\tau'$, $\langle \emptyset \rangle \cdot u$ is consistent with $\langle \alpha_n \rangle \tau'$, so by Lem. 26, the assignment ν s.t. $\nu(v_0) = \alpha_n(v)$ and $\nu(v) = \alpha'_k(v)$ for all $v \in V$ satisfies $h(u)$. Thus Lem. 29 implies that $\text{CG}_\psi(P)$ has a path $(P, \varphi_{init}) \rightarrow_u^* (P', \varphi)$ s.t. $\varphi \sim h(u)$. As $\nu \models h(u)$, by Def. 28, $\alpha_n \models (\exists \varphi)(\bar{V})$ because $\nu(v_0) = \alpha_n(v)$ for all $v \in V$. As P' is non-final, it follows that $\alpha_n \models \text{FUNS}(G)$, a contradiction; so τ' cannot exist. Thus $\tau \models \llbracket \psi = \text{PS} \rrbracket$. The case where $\tau \not\models \psi$ is dual, using $\text{FSAT}(G)$. \square

Finally, if ψ has a finite summary (Φ, \sim) , by Lem. 29 a constraint graph $\text{CG}(P, \sim)$ computed from \mathcal{D}_ψ can take all nodes from the finitely many equivalence classes of Φ , so that the constraint graph can be finite. Therefore:

Corollary 31. Monitoring is solvable for $\psi \in \mathcal{L}$ if ψ has finite summary.

Concrete criteria for solvability

We next apply Cor. 31 to show solvability of concrete property classes, by restricting the constraint set or control flow.

Monotonicity constraints (MCs) over variables \mathcal{V} and domain D have the form $p \odot q$ where $p, q \in D \cup \mathcal{V}$ and \odot is one of $=, \neq, \leq$, or $<$. We call an LTL_f property whose constraint atoms are MCs over D an *MC_D property*. Ex. 11 gives a simple example of an MC property. However, MC properties are rich enough to capture practically important settings: For instance, Geist, Rozier, and Schumann (2014) model the specification of a fluxgate magnetometer of a

Swift UAS system operated by NASA, as linear-time properties with arithmetic comparisons of sensor values, which are all MCs over \mathbb{Q} (cf. their Table 2). We have that:

Theorem 32. *Monitoring is solvable for $MC_{\mathbb{Q}}$ properties.*

Proof (sketch). This is shown as in (Felli, Montali, and Winkler 2022b, Thm. 5.2), exploiting that quantifier elimination of MC formulas over \mathbb{Q} produces MC formulas with the same constants, so that a finite history set exists. \square

Integer periodicity constraints (IPCs) confine the constraint language in a similar way as MCs, but allow equality modulo, and variable-to-variable comparisons are restricted. IPCs are e.g. used in calendar formalisms (Demri 2006). More precisely, IPC atoms have the form $x = y$, $x \odot d$ for $\odot \in \{=, \neq, <, >\}$, $x \equiv_k y + d$, or $x \equiv_k d$, for variables x, y with domain \mathbb{Z} and $k, d \in \mathbb{N}$. An \mathcal{L} property over IPC atoms is called an *IPC property*. The next result is proven similarly as Thm. 36, cf. (Felli, Montali, and Winkler 2022a, Thm. 4).

Theorem 33. *Monitoring is solvable for IPC properties.*

A simple example of an IPC property is $(x \equiv_7 y + 1) \cup (x = z)$. Also e.g. a parallel program as in (Havelund, Reger, and Rosu 2019, Fig. 3) can be modeled as an IPC property, where monitoring detect at runtime whether a trace may trigger the race on variable x that the program exhibits.

Control flow restrictions can be used instead of confining the constraint language to obtain solvability. *Computation graphs* (Damaggio, Deutsch, and Vianu 2012) were introduced to track dependencies between variable instances at different trace events. A property ψ is said to have *k-bounded lookback* (Felli, Montali, and Winkler 2022b) if all paths in computation graphs of ψ have length at most k , without counting equality edges. Intuitively, this corresponds to forbid variable updates that depend on an unbounded history of values (of the same or of other variables), but only on a k -bounded “moving window”. A property has bounded lookback (BL) if its has k -bounded lookback for some $k \geq 0$. E.g., $G(x' > x)$ does not have BL, because there is an unbounded dependency chain between the values of x , whereas $F(x' > 2y) \wedge G(x + y > 0)$ has 1-bounded lookback as variable comparisons span at most one time unit. Also all properties with lookahead 0 have BL. Monitoring of BL properties is solvable because, roughly, the quantifier depth of history constraints is upper-bounded a priori. A formal definition, proof and examples are in (Felli et al. 2022).

Theorem 34. *Monitoring of BL properties is solvable.*

Gap-order constraints (GCs) have the form $x - y \geq k$ for x and y either variables with domain \mathbb{Z} or integers, and $k \in \mathbb{N}$. Satisfiability of GC properties is decidable (Bozzelli and Pinchinat 2014). However, validity is undecidable, because a GC property ψ is valid iff $\neg\psi$ is unsatisfiable, where $\neg\psi$ can be written as a property whose atoms are negated GCs, i.e. of the form $\neg c$ for c a GC. Negated GCs can model counters (Bozzelli and Pinchinat 2014, Thm. 12), so that state reachability of a 2-counter machine can be reduced to satisfiability of a property over negated GCs. Consequently,

Remark 35. *Monitoring of GC properties is not solvable.*

MCs over \mathbb{Z} can be written as GCs (Bozzelli and Pinchinat 2014), with the key difference that for an MC c also $\neg c$ is an MC, and hence a GC. Thus, we can assume that in a DFA for an $MC_{\mathbb{Z}}$ property ψ , all constraints are GCs.

In (Felli, Montali, and Winkler 2022b, Thm. 5.5) it was shown that a transition system over GCs admits a finite summary (GC_K, \sim_K) . Here, GC_K is the finite set of quantifier-free formulas with GC atoms over variables V and constants $\leq K$, and \sim_K is the cutoff equivalence relation, where equivalence of two formulas is checked after replacing constants greater than K by K , for K the maximal difference between constants in the input. Finite summary relies on the initial assignment of V being fixed, whereas by Def. 25 history constraints start from φ_{init} , i.e., a parametric assignment $\bar{V} \mapsto \bar{V}_0$. This can hamper finiteness of the summary. E.g., for $\psi = X_w(x_{cur} > x_{pre} \cup x_{cur} > 5)$ there are finitely many equivalence classes only if the initial value z of x is known (after at most $5 - z$ steps ψ holds). Nonetheless, Cor. 31 can be used to show that monitoring is solvable, although the procedure needs to be suitably refined.

Theorem 36. *Monitoring is solvable for $MC_{\mathbb{Z}}$ properties.*

Proof (sketch). We modify our procedure as follows: (1) in line 5, we compute the CG starting from $\bigwedge v = \alpha_n(v)$ instead of φ_{init} . By (Felli, Montali, and Winkler 2022b, Thm. 5.5) a finite summary exists, so the CG computation terminates. (2) The checks in lines 8,9 are simplified: as the initial assignment is now integrated in the CG, at line 8 it suffices to check whether the CG has a non-empty path to a non-final state (resp. to a final state in line 9). It is not hard to adapt the proof of Thm. 30 accordingly. However, the changed procedure comes with the price of computing CGs not only for every DFA state but also every assignment in the trace. \square

Conclusions

Implementation. We implemented our approach in a prototype, whose source code and web interface are available via <https://bit.ly/3QFoJHA>. The tool takes an ALTL_f property ψ and a trace as input, determines whether ψ is in one of the decidable classes identified in the last section, constructs and visualizes a monitor, and computes the monitoring state. The tool is implemented in Python, using Z3 (de Moura and Bjørner 2008) and CVC5 (Detert et al. 2014) for SMT checks and quantifier elimination.

Future work. We want to lift our automata-based approach to the case of traces of richer states, equipped with full-fledged relations. Hence we plan to study how to integrate our approach with (Decker, Leucker, and Thoma 2016) and (Calvanese et al. 2022). The former considers states consisting of first-order structures modulo theories, but does not foresee any form of lookahead in the properties. In (Calvanese et al. 2022), instead, states are labeled by first-order interpretations and properties are expressed in a fragment of first-order LTL with a controlled first-order quantification across time; however arithmetic theories are not supported.

Acknowledgments This work was partially funded by the Unibz project SMART-APP and the PRIN MIUR project PINPOINT.

References

- Bartocci, E.; Falcone, Y.; Francalanza, A.; and Reger, G. 2018. Introduction to Runtime Verification. In Bartocci, E.; and Falcone, Y., eds., *Lectures on Runtime Verification - Introductory and Advanced Topics*, volume 10457 of *LNCS*, 1–33. Springer.
- Basin, D. A.; Klaedtke, F.; Müller, S.; and Zalinescu, E. 2015. Monitoring Metric First-Order Temporal Properties. *J. ACM*, 62(2): 15:1–15:45.
- Bauer, A.; Leucker, M.; and Schallhart, C. 2010. Comparing LTL Semantics for Runtime Verification. *J. Logic and Comput.*, 20(3): 651–674.
- Bozzelli, L.; and Pinchinat, S. 2014. Verification of gap-order constraint abstractions of counter systems. *Theor. Comput. Sci.*, 523: 1–36.
- Calvanese, D.; De Giacomo, G.; Montali, M.; and Patrizi, F. 2022. Verification and Monitoring for First-Order LTL with Persistence-Preserving Quantification over Finite and Infinite Traces. In *Proc. 31st IJCAI*, 2553–2560.
- Damaggio, E.; Deutsch, A.; and Vianu, V. 2012. Artifact systems with data dependencies and arithmetic. *ACM Trans. Database Syst.*, 37(3): 22:1–22:36.
- D’Angelo, B.; Sankaranarayanan, S.; Sánchez, C.; Robinson, W.; Finkbeiner, B.; Sipma, H. B.; Mehrotra, S.; and Manna, Z. 2005. LOLA: Runtime Monitoring of Synchronous Systems. In *Proc. 12th TIME*, 166–174.
- Dastani, M.; Torroni, P.; and Yorke-Smith, N. 2018. Monitoring norms: a multi-disciplinary perspective. *Knowl. Eng. Rev.*, 33.
- De Giacomo, G.; De Masellis, R.; Maggi, F. M.; and Montali, M. 2022. Monitoring Constraints and Metaconstraints with Temporal Logics on Finite Traces. *ACM Trans. Softw. Eng. Methodol.*, 31(4).
- De Giacomo, G.; and Vardi, M. Y. 2013. Linear Temporal Logic and Linear Dynamic Logic on Finite Traces. In Rossi, F., ed., *Proc. 23rd IJCAI*, 854–860. IJCAI/AAAI.
- de Moura, L.; and Bjørner, N. 2008. Z3: An Efficient SMT Solver. In *Proc. 14th TACAS*, volume 4963 of *LNCS*, 337–340.
- Decker, N.; Leucker, M.; and Thoma, D. 2016. Monitoring modulo theories. *Int. J. Softw. Tools Technol. Transf.*, 18(2): 205–225.
- Demri, S. 2006. LTL over integer periodicity constraints. *Theor. Comput. Sci.*, 360(1-3): 96–123.
- Demri, S.; and D’Souza, D. 2007. An automata-theoretic approach to constraint LTL. *Inform. Comput.*, 205(3): 380–415.
- Deters, M.; Reynolds, A.; King, T.; Barrett, C. W.; and Tinelli, C. 2014. A tour of CVC4: How it works, and how to use it. In *Proc. 14th FMCAD*, 7.
- Falcone, Y.; Krstic, S.; Reger, G.; and Traytel, D. 2021. A taxonomy for classifying runtime verification tools. *Int. J. Softw. Tools Technol. Transf.*, 23(2): 255–284.
- Faymonville, P.; Finkbeiner, B.; Schledjewski, M.; Schwenger, M.; Stenger, M.; Tentrup, L.; and Torfah, H. 2019. StreamLAB: Stream-based Monitoring of Cyber-Physical Systems. In *Proc. 31st CAV*, volume 11561 of *LNCS*, 421–431.
- Felli, P.; Montali, M.; Patrizi, F.; and Winkler, S. 2022. Monitoring Arithmetic Temporal Properties on Finite Traces (Extended Version). Available from <https://arxiv.org/abs/2211.17166>.
- Felli, P.; Montali, M.; and Winkler, S. 2022a. CTL* model checking for data-aware dynamic systems with arithmetic. In *Proc. 11th IJCAR*, volume 13385, 36–56.
- Felli, P.; Montali, M.; and Winkler, S. 2022b. Linear-Time Verification of Data-Aware Dynamic Systems with Arithmetic. In *Proc. 36th AAI*, 5642–5650. AAI Press. Extended version available from <https://arxiv.org/abs/2203.07982>.
- Geatti, L.; Gianola, A.; and Gigante, N. 2022. Linear Temporal Logic Modulo Theories Over Finite Traces. In *Proc. 31st IJCAI*, 2641–2647.
- Geist, J.; Rozier, K. Y.; and Schumann, J. 2014. Runtime Observer Pairs and Bayesian Network Reasoners On-board FPGAs: Flight-Certifiable System Health Management for Embedded Systems. In *Proc. 5th Runtime Verification*, volume 8734 of *LNCS*, 215–230.
- Havelund, K.; Reger, G.; and Rosu, G. 2019. Runtime Verification Past Experiences and Future Projections. In *Computing and Software Science - State of the Art and Perspectives*, volume 10000 of *LNCS*, 532–562. Springer.
- Leucker, M.; and Schallhart, C. 2009. A brief account of runtime verification. *J. Log. Algebraic Methods Program.*, 78(5): 293–303.
- Ly, L. T.; Maggi, F. M.; Montali, M.; Rinderle-Ma, S.; and van der Aalst, W. M. P. 2015. Compliance monitoring in business processes: Functionalities, application, and tool-support. *Inf. Syst.*, 54: 209–234.
- Maggi, F. M.; Francescomarino, C. D.; Dumas, M.; and Ghidini, C. 2014. Predictive Monitoring of Business Processes. In *Proc. 26th CAiSE*, volume 8484 of *LNCS*, 457–472. Springer.
- Maggi, F. M.; Montali, M.; Westergaard, M.; and van der Aalst, W. M. P. 2011. Monitoring Business Constraints with Linear Temporal Logic: An Approach Based on Colored Automata. In *Proc. 9th BPM*, volume 6896 of *LNCS*, 132–147. Springer.
- Presburger, M. 1929. Über die Vollständigkeit eines gewissen Systems der Arithmetik ganzer Zahlen, in welchem die Addition als einzige Operation hervortritt. In *Comptes Rendus du I congrès de Mathem. des Pays Slaves*, 92–101.
- Reger, G.; Cruz, H. C.; and Rydeheard, D. E. 2015. MarQ: Monitoring at Runtime with QEA. In *Proc. 21st TACAS*, volume 9035 of *LNCS*, 596–610. Springer.
- Xu, H.; and Cheng, Y. 2007. Model checking bidding behaviors in internet concurrent auctions. *Comput. Syst. Sci. Eng.*, 22(4).