Runtime Verification: From Propositional to First-Order Temporal Logic^{*}

Klaus Havelund¹ and Doron Peled²

¹Jet Propulsion Laboratory, California Institute of Technology, USA ² Department of Computer Science Bar Ilan University, Israel

Abstract. Runtime Verification is a branch of formal methods concerned with analysis of execution traces for the purpose of determining the state or general quality of the executing system. The field covers numerous approaches, one of which is specification-based runtime verification, where execution traces are checked against formal specifications. The paper presents syntax, semantics, and monitoring algorithms for respectively propositional and first-order temporal logics. In propositional logics the observed events in the execution trace are represented using atomic propositions, while first-order logic allows universal and existential quantification over data occurring as arguments in events. Monitoring of the first-order case is drastically more challenging than the propositional case, and we present a solution for this problem based on BDDs. We furthermore discuss monitorability of temporal properties by dividing them into different classes representing different degrees of monitorability.

1 Introduction

Runtime verification (RV) [2, 16] allows monitoring (analysis) of executions of a system, directly, without the need for modeling the system. It has some commonality with other formal methods such as testing, model checking and formal verification, including the use of a specification formalisms¹. However, it differs a lot in goals, the algorithms used, and the complexity and the coverage it suggests. Model checking performs a comprehensive search on a model of the system under test. Testing generates inputs to drive system executions, trying to provide a good coverage, yet keeping the complexity low, at the price of losing exhaustiveness. Formal verification attempts full proof of correctness based on deductive techniques. Runtime verification does not directly concern

^{*} The research performed by the first author was carried out at Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. The research performed by the second author was partially funded by Israeli Science Foundation grant 2239/15: "Runtime Measuring and Checking of Cyber Physical Systems".

¹ RV can be understood more broadly to mean: any processing of execution traces for the purpose of evaluating a system state or quality. Some approaches do not involve specifications but rather use pre-programmed algorithms as monitors.

itself with coverage and the selection of execution paths, but rather focuses on analyzing a single execution trace (or a collection thereof). An execution trace is generated by the observed executing system, typically by instrumenting the system to generate events as important transitions take place. Instrumentation can be manual by inserting logging statements in the code, or it can be automated using instrumentation software, such as aspect-oriented programming frameworks.

Runtime verification can take place on-line, as the system executes, or off-line, by processing log files produced by the system. In the case of on-line processing, runtime verification obtains the information about the execution as it unfolds, oftentimes without seeing the complete sequence; yet it is required to provide a verdict as soon as possible. The critical complexity measure here is the *incremental complexity*, which is performed for each new event reported to the monitor. The calculation needs to be fast enough to keep in pace with the executing system.

Following in part [26] and [17], we present algorithms for the runtime verification of linear temporal logic properties, which is the most common specification formalism used for both runtime verification and model checking. We start with the propositional case, where an execution trace is checked against a future or past time propositional LTL formula. For an online algorithm, which observes the execution trace event by event, a verdict is not guaranteed in any finite time. Runtime monitorability identifies what kind of verdicts can be expected when monitoring an execution against a temporal property. Monitoring temporal properties is often restricted to safety properties. There are two main reasons for this restriction: the first is that the algorithm for checking safety is rather efficient, polynomial in the size of the property; the other reason is that for safety properties we are guaranteed to have a finite evidence for a negative verdict (albeit there is not always a bound on when such an evidence can be given).

After presenting the theory of monitoring propositional temporal logic we move on to the more demanding challenge of monitoring properties that depend on data reported to the monitor. This can be handled by a parametrized version of temporal logic (or using parametrized automata), but more generally it calls for using a first-order version of the temporal logic. We will concentrate on first-order safety properties. One of the challenges here is that the data may, in principle, be unbounded and we only learn about the actual values that are monitored as they appear in reported events. Another problem is that, unlike the propositional case, the amount of data that needs to be kept may keep growing during the execution. This calls for a clever representation that allows fast processing of many data elements. We present an algorithm based on BDDs, which is implemented in the tool DEJAVU [15].

The paper is organized as follows. Section 2 introduces propositional linear temporal logic, including future as well as past time operators, its syntax, semantics, and some pragmatics. Section 3 presents a general theory of monitorability of temporal properties, such as those formulated in LTL. Section 4 outlines algorithms for monitoring propositional LTL properties, first future time, and then past time. Section 5 introduces first-order past time LTL, its syntax, semantics, and some pragmatics. Section 6 outlines an algorithm for monitoring first-order past LTL properties. Finally, Section 7 concludes the paper.

2 Propositional LTL

The definition of linear temporal logic including future and past time operators is as follows [23]:

$$\varphi ::= true |p| (\phi \land \phi) |\neg \phi| (\phi \mathcal{U} \phi) | \bigcirc \phi | (\phi \mathcal{S} \phi) | \ominus \phi$$

where *p* is a proposition from a finite set of propositions *P*, with \mathcal{U} standing for *until*, \bigcirc standing for *next-time*, \mathcal{S} standing for *since*, and \ominus standing for *previous-time*. One can also write $(\phi \lor \psi)$ instead of $\neg(\neg \phi \land \neg \psi)$, $(\phi \rightarrow \psi)$ instead of $(\neg \phi \lor \psi)$, $\diamond \phi$ (*eventually* ϕ) instead of $(true \mathcal{U} \phi)$, $\Box \phi$ (*always* ϕ) instead of $\neg \diamond \neg \phi$, **P** ϕ (*past* ϕ) instead of (*true* $\mathcal{S} \phi$) and **H** ϕ (*history* ϕ) instead of $\neg \mathbf{P} \neg \phi$.

LTL formulas are interpreted over an infinite sequence of events $\xi = e_1.e_2.e_3...$, where $e_i \subseteq P$ for each i > 0. These are the propositions that *hold* in that event. LTL's semantics is defined as follows:

- $-\xi$, $i \models true$.
- ξ , $i \models p$ iff $p \in e_i$.
- $-\xi$, $i \models \neg \varphi$ iff not ξ , $i \models \varphi$.
- $\xi, i \models (\phi \land \psi)$ iff $\xi, i \models \phi$ and $\xi, i \models \psi$.
- $-\xi, i \models \bigcirc \varphi \text{ iff } \xi, i+1 \models \varphi.$
- ξ , $i \models (\varphi \ U \psi)$ iff for some $j \ge i, \xi, j \models \psi$, and for all $i \le k < j$ it holds that $\xi, k \models \varphi$.
- ξ , *i* $\models \ominus \varphi$ iff *i* > 1 and ξ , *i* 1 $\models \varphi$.
- $\xi, i \models (\varphi S \Psi)$ iff $\xi, i \models \Psi$ or the following hold²: i > 1, $\xi, i \models \varphi$ and $\xi, i 1 \models (\varphi S \Psi)$.

Then $\xi \models \varphi$ when ξ , $1 \models \varphi$.

This definition of propositional temporal logic contains both future modalities (\mathcal{U} , \Box , \diamond and \bigcirc) and past modalities (\mathcal{S} , **H**, **P** and \ominus). However, we do not always need to use all off them:

- Removing the past temporal operators does not affect the expressiveness of the logic [13]. On the other hand, there are examples of properties that are much more compact when expressed using both the past and the present operators.
- Properties of the form $\Box \varphi$, where φ does *not* contain the future operators form an important class. There are several reasons for restricting runtime verification to such properties. These properties correspond to temporal *safety properties* [22, 1]: failure can always be detected on a finite prefix [7]. Moreover, expressing safety properties in this form allows an efficient runtime verification algorithm that is only polynomial in the size of the specification [17]³.

² This definition is equivalent to the traditional definition $\xi, i \models (\varphi \ S \ \psi)$ iff for some $0 < j \le i$, $\xi, j \models \psi$, and for all $j < k \le i$ it holds that $\xi, k \models \varphi$, but is more intuitive for the forthcoming presentation of the RV algorithm.

³ There are examples of safety properties that are much more compact when expressed with the past temporal operators [24], and for symmetrical considerations also vice versa.

3 Monitorability of Propositional Temporal Logic

Online runtime verification of LTL properties inspects finite prefixes of the execution. Assume an observed system *S*, and assume further that a finite execution of *S* up to a certain point is captured as an execution trace $\xi = e_1.e_2. \ldots .e_n$, which is a sequence of observed events, each of type \mathbb{E} . Each event e_i captures a snapshot of *S*'s execution state. Then the RV problem can be formulated as constructing a program *M* with the type $M : \mathbb{E}^* \to D$, which when applied to the trace ξ , as in $M(\xi)$, returns some data value $d \in D$ in a domain *D* of interest. Typically *M* is generated from a formal specification, given as a temporal logic formula or a state machine. Because online RV observes at each time only a finite part of the execution, it can sometimes provide only a partial verdict on the satisfaction and violation of the inspected property [6, 25]. This motivates providing three kinds of verdicts as possible values for the domain *D*:

- *failed* when the current prefix cannot be extended in any way into an execution that satisfies the specification,
- *satisfied* when any possible extension of the current prefix satisfies the specification, and
- *undecided* when the current prefix can be extended to satisfy the specification but also extended to satisfy its negation.

For example, the property $\Box p$ (for some atomic proposition p), which asserts that p always happens, can be refuted by a runtime monitor if p does not hold in some observed event. At this point, no matter which way the execution is extended, the property will not hold, resulting in a *failed* verdict. However, no finite prefix of an execution can establish that $\Box p$ holds. In a similar way, the property $\Diamond p$ cannot be refuted, since p may appear at any time in the future; but once p happens, we know that the property is satisfied, independent on any continuation, and we can issue a *satisfied* verdict. For the property $(\Box p \lor \Diamond q)$ we may not have a verdict at any finite time, in the case where all the observed events satisfy both p and $\neg q$. On the other hand, we may never "lose hope" to have such a verdict, as a later event satisfying q will result in a positive verdict: at this point we can abandon the monitoring, since the property cannot be further violated. On the other hand, for the property $\Box \diamond p$ we can never provide a verdict in finite time: for whatever happens, p can still appear an infinite number of times, and we cannot guarantee or refute that this property holds when observing any finite prefix of an execution. The problem of monitorability of a temporal property was studied in [7, 12, 27], basically requiring that at any point of monitoring we still have a possibility to obtain a finite positive or negative verdict.

Safety and liveness temporal properties were defined informally on infinite execution sequences by Lamport [22] as *something bad cannot happen* and *something good will happen*. These informal definitions were later formalized by Alpern and Schneider [1]. Guarantee properties where used in an orthogonal characterization by Manna and Pnueli [23]. Guarantee properties are the dual of safety properties, that is, the negation of a safety property is a guarantee property and vice versa.

These classes of properties can be seen as characterizing finite monitorability of temporal properties: if a safety property is violated, there will be a finite prefix witnessing it; on the other hand, for a liveness property, one can never provide such a finite negative evidence. We suggest the following alternative definitions of classes of temporal properties.

- **AFR / safety** Always Finitely Refutable: when the property does not hold on an infinite execution, a *failed* verdict can be identified after a finite prefix.
- **AFS / guarantee** Always Finitely Satisfiable: when the property is satisfied on an infinite execution, a *satisfied* verdict can be identified after a finite prefix.
- **NFR / liveness** Never Finitely Refutable: when the property does not hold on an infinite execution, refutation can never be identified after a finite prefix.
- **NFS / morbidity** Never Finitely Satisfiable: When the property is satisfied on an infinite execution, satisfaction can never be identified after a finite prefix.

It is easy to see that the definitions of the classes AFR and safety in [1] are the same and so are those for AFS and guarantee. A liveness property φ is defined to satisfy that any finite prefix can be extended to an execution that satisfies φ . The definition of the class NFR only mentions prefixes of executions that do not satisfy φ ; but for prefixes of executions that satisfy φ this trivially holds. The correspondence between NFS and morbidity is shown in a symmetric way.

The above four classes of properties, however, do not cover the entire set of possible temporal properties, independent of the actual formalism that is used to express them. The following two classes complete the classification.

- **SFR** Sometimes Finitely Refutable: for some infinite executions that violate the property, refutation can be identified after a finite prefix; for other infinite executions violating the property, this is not the case.
- **SFS** Sometimes Finitely Satisfiable: for some infinite executions that satisfy the property, satisfaction can be identified after a finite prefix; for other infinite executions satisfying the property, this is not the case.

Bauer, Leucker and Schallhart [7] define three categories of prefixes of elements from 2^{P} .

- A *good* prefix is one where all its extensions (infinite sequences of elements from 2^{P}) satisfy the monitored property φ .
- A *bad* prefix is one where none of its infinite extensions satisfy φ .
- An *ugly* prefix cannot be extended into a good or a bad prefix.

When identifying a good or a bad finite prefix, we are done tracing the execution and can announce that the monitored property is satisfied or failed, respectively. After an ugly prefix, satisfaction or refutation of φ depends on the entire infinite execution, and cannot be determined in finite time. Note that a property has a good prefix if it is *not* a morbidity property, and a bad prefix if it is *not* a liveness property. *Monitorability* of a property φ is defined in [7] as the lack of ugly prefixes for the property φ . This definition is consistent with [27].

Any property that is in AFR (safety) or in AFS (guarantee) is monitorable [7, 12]. A property that is NFR \cap NFS is non-monitorable. In fact no verdict is ever expected on any sequence that is monitored against such a property. This leaves the three classes SFR \cap SFS, SFR \cap NFS and NFR \cap SFS, for which some properties are monitorable and others are not. This is demonstrated in the following table.



Fig. 1: Classification of properties: safety, guarantee, liveness, morbidity and quaestio.

Class	monitorable example	non-monitorable example
$SFR \cap SFS$	$((\Diamond r \lor \Box \Diamond p) \land \bigcirc q)$	$((p \lor \Box \diamondsuit p) \land \bigcirc q)$
$SFR \cap NFS$	$(\Diamond p \land \Box q)$	$(\Box \diamondsuit p \land \bigcirc q)$
$NFR \cap SFS$	$(\Box p \lor \Diamond q)$	$(\Box \diamondsuit p \lor \bigcirc q)$

The set of all properties *Prop* is not covered by safety, guarantee, liveness and morbidity. The missing properties are in SFR \cap SFS. We call the class of such properties **Quaestio** (Latin for *question*). Figure 1 presents the relationship between the different classes of properties and their intersections, with LTL specification examples.

4 Runtime Verification for Propositional LTL

4.1 Runtime Verification for Propositional Future LTL

We present three algorithms. The first one is a classical algorithm for runtime verification of LTL (or Büchi automata) properties. The second algorithm can be used to check during run time what kind of verdicts can still be produced given the current prefix. The third algorithm can be used to check whether the property is monitorable.

Algorithm 1: Monitoring sequences using automata

Kupferman and Vardi [20] provide an algorithm for detecting good and bad prefixes. For good prefixes, start by constructing a Büchi automaton $\mathcal{A}_{\neg\phi}$ for $\neg\phi$, e.g., using the translation in [14]. Note that this automaton is not necessarily deterministic [30]. States of $\mathcal{A}_{\neg\varphi}$, from which one cannot reach a cycle that contains an accepting state, are deleted. Checking for a positive verdict for φ , one keeps for each monitored prefix the set of states that $\mathcal{A}_{\neg\varphi}$ would be in after observing that input. One starts with the set of initial states of the automaton $\mathcal{A}_{\neg\varphi}$. Given the current set of successors *S* and an event $e \in 2^P$, the next set of successors *S'* is set to the successors of the states in *S* according to the transition relation Δ of $\mathcal{A}_{\neg\varphi}$. That is, $S' = \{s' | s \in S \land (s, e, s') \in \Delta\}$. Reaching the empty set of states, the monitored sequence is good, and the property must hold since the current prefix cannot be completed into an infinite execution satisfying $\neg \varphi$.

This is basically a *subset construction* for a deterministic automaton \mathcal{B}_{φ} , whose initial state is the set of initial states of $\mathcal{A}_{\neg\varphi}$, accepting state is the empty set, and transition relation as described above. The size of this automaton is $O(2^{2^{|P|}})$, resulting in double exponential explosion from the size of the checked LTL property. But in fact, we do not need to construct the entire automaton \mathcal{B}_{φ} in advance, and can avoid the double exponential explosion by calculating its current state on-the-fly, while performing runtime verification. Thus, the incremental processing per each event is exponential in the size of the checked LTL property. Unfortunately, a single exponential explosion is unavoidable [20].

Checking for a *failed* verdict for φ is done with a symmetric construction, translating φ into a Büchi automaton \mathcal{A}_{φ} and then the deterministic automaton $\mathcal{B}_{\neg\varphi}$ (or calculating its states on-the-fly) using a subset construction as above. Note that $\mathcal{A}_{\neg\varphi}$ is used to construct \mathcal{B}_{φ} and \mathcal{A}_{φ} is used to construct $\mathcal{B}_{\neg\varphi}$. Runtime verification of φ uses both automata for the monitored input, reporting a *failed* verdict if $\mathcal{B}_{\neg\varphi}$ reaches an accepting state, a *satisfied* verdict if \mathcal{B}_{φ} reaches an accepting state, and an *undecided* verdict otherwise. The algorithm guarantees to report a positive or negative verdict on the minimal good or bad prefix that is observed.

Algorithm 2: Checking availability of future verdicts

We alter the above runtime verification algorithm to check whether positive or negative verdicts can still be obtained after the current monitored prefix at runtime. Applying DFS on \mathcal{B}_{ϕ} , we search for states from which one cannot reach the accepting state. We replace these states with a single state \perp with a self loop, obtaining the automaton \mathcal{C}_{ϕ} . Reaching \perp , after monitoring a finite prefix σ with \mathcal{C}_{ϕ} happens exactly when we will not have a good prefix anymore. This means that after σ , a *satisfied* verdict cannot be issued anymore for ϕ .

Similarly, we perform BFS on $\mathcal{B}_{\neg \varphi}$ to find all the states in which the accepting state is not reachable, then replace them by a single state \top with a self loop, obtaining $\mathcal{C}_{\neg \varphi}$. Reaching \top after monitoring a prefix means that we will not be able again to have a bad prefix, hence a *failed* verdict cannot be issued anymore for φ . There is no point in continuing the monitoring if we reach the pair of states (\bot, \top) , since there is no further information, positive or negative, that will be later given by the runtime verification. This happens when the currently monitored prefix is ugly.

We can perform runtime verification while updating the state of both automata, C_{ϕ} and $C_{\neg\phi}$ on-the-fly, upon each input event. However, we need to be able to predict if,

from the current state, an accepting state is not reachable. While this can be done in space polynomial in the size of C_{ϕ} and $C_{\neg\phi}$, it makes an incremental calculation whose time complexity is doubly exponential in the size of ϕ , as is the algorithm for that by Pnueli and Zaks [27]. This is hardly a reasonable complexity for the incremental calculation performed between successive monitored events for an on-line algorithm. Hence, a pre-calculation of these two automata before the monitoring starts is preferable, leaving the incremental complexity exponential in ϕ , as in Algorithm 1.

Algorithm 3: Checking monitorability

A small variant on the construction of C_{φ} and $C_{\neg\varphi}$ allows checking if a property is monitorable. The algorithm is simple: construct the product $C_{\varphi} \times C_{\neg\varphi}$ and check whether the state (\bot, \top) is reachable. If so, the property is non-monitorable, since there is a prefix that will transfer the product automaton to this state and thus it is ugly. It is not sufficient to check separately that C_{φ} can reach \top and that $C_{\neg\varphi}$ can reach \bot . In the property $(\neg(p \wedge r) \wedge ((\neg p \mathcal{U}(r \wedge \Diamond q)) \vee (\neg r \mathcal{U}(p \wedge \Box q))))$: both \bot and \top can be reached, separately, depending on which of the predicates r or p happens first. But in either case, there is still a possibility for a good or a bad extension, hence it is a monitorable property.

If the automaton $C_{\phi} \times C_{\neg \phi}$ consists of only a single state (\bot, \top) , then there is no information whatsoever that we can obtain from monitoring the property.

The above algorithm is simple enough to construct, however its complexity is doubly exponential in the size of the given LTL property. This may not be a problem, as the algorithm is performed off-line and the LTL specifications are often quite short.

We show that checking monitorability is in EXPSPACE-complete. The upper bound is achieved by a binary search version of this algorithm⁴. For the lower bound we show a reduction from checking if a property is (not) a liveness property, a problem known to be in EXPSPACE-complete [28, 21].

- We first neutralize bad prefixes. Now, when ψ is satisfiable, then $\Diamond \psi$ is monitorable (specifically, any prefix can be completed into a *good* prefix) iff ψ has a good prefix.
- Checking satisfiability of a property ψ is in PSPACE-complete [29]⁵.
- ψ has a good prefix iff ψ is not a morbidity property, i.e., if $\phi = \neg \psi$ is not a liveness property.
- Now, φ is *not* a liveness property iff either φ is valid or $\Diamond \neg \varphi$ is monitorable.

4.2 Runtime Verification for Propositional Past LTL

Algorithm

The algorithm for past LTL, first presented in [17], is based on the observation that the semantics of the past time formulas $\ominus \varphi$ and $(\varphi S \psi)$ in the current step *i* is defined in terms of the semantics in the previous step *i* – 1 of a subformula, here recalled from Section 2:

⁴ To show that a property is not monitorable, one needs to guess a state of $\mathcal{B}_{\phi} \times \mathcal{B}_{\neg\phi}$ and check that (1) it is reachable, and (2) one cannot reach from it an empty component, both for \mathcal{B}_{ϕ} and for $\mathcal{B}_{\neg\phi}$. (There is no need to construct \mathcal{C}_{ϕ} or $\mathcal{C}_{\neg\phi}$.)

⁵ Proving that liveness was PSPACE-hard was shown in [3].

- $\xi, i \models \ominus \varphi$ iff i > 1 and $\xi, i - 1 \models \varphi$. - $\xi, i \models (\varphi S \psi)$ iff $\xi, i \models \psi$ or the following hold: $i > 1, \xi, i \models \varphi$ and $\xi, i - 1 \models (\varphi S \psi)$.

One only needs to look one step, or event, backwards in order to compute the new truth value of a formula and of its subformulas. The algorithm, shown below, operates on two vectors (arrays) of values indexed by subformulas: pre for the state before that event, and now for the current state (after the last seen event).

- 1. Initially, for each subformula φ , now(φ) := *false*.
- 2. Observe a new event (as a set of ground predicates) s as input.
- 3. Let pre := now.
- 4. Make the following updates for each subformula. If ϕ is a subformula of ψ then $\mathsf{now}(\phi)$ is updated before $\mathsf{now}(\psi)$.
 - now(true) := true.
 - $\mathsf{now}((\phi \land \psi)) := \mathsf{now}(\phi) \text{ and } \mathsf{now}(\psi).$
 - now($\neg \phi$) := *not* now(ϕ).
 - $\mathsf{now}((\varphi \ S \ \psi)) := \mathsf{now}(\psi) \ or \ (\mathsf{now}(\varphi) \ and \ \mathsf{pre}((\varphi \ S \psi))).$
 - now($\ominus \phi$) := pre(ϕ).
- 5. Goto step 2.

An Example

As an example⁶, consider the formula *close* $\land \ominus open$, and suppose we evaluate it against the trace *open.close* at step i = 2 (after seeing the *close* event). The algorithm performs the following assignments, resulting in the formula becoming true (assuming that pre(*open*) is true):

$$\begin{array}{l} \mathsf{now}(open) := false\\ \mathsf{now}(close) := true\\ \mathsf{now}(\ominus open) := \mathsf{pre}(open)\\ \mathsf{now}(close \land \ominus open) := \mathsf{now}(close) \land \mathsf{now}(\ominus open) \end{array}$$

The above suggested algorithm *interprets* a formula on a trace. As an alternative we can *synthesize* a program that is specialized for monitoring the property as in [17]. Figure 2 (left) shows a generated monitor program for the property. Two Boolean valued arrays pre for the previous state and now are declared and operated on. The indices 0...3 correspond to the enumeration of the subformulas shown in the Abstract Syntax Tree (AST) in Figure 2 (right). For each observed event, the function evaluate() computes the now array from highest to lowest index, and returns true (property is satisfied in this position of the trace) iff now(0) is *true*.

⁶ All examples of safety properties henceforth will omit the implied \Box operator.



Fig. 2: Monitor (left) and AST (right) for propositional property.

5 First-Order Past LTL

First-order past LTL allows universal and existential quantification over data occurring as parameters in events. Such events are referred to as *predicates* (or *parametric events*). Consider a predicate open(f), indicating that a file f is being opened, and a predicate close(f) indicating that f is being closed. We can formulate that a file cannot be closed unless it was opened before with the following first-order past LTL formula:

$$\forall f(close(f) \longrightarrow \mathbf{P}open(f)) \tag{1}$$

Here **P** is the "sometimes in the past" temporal operator. This property must be checked for every monitored event. Already in this very simple example we see that we need to store *all* the names of files that were previously opened so we can compare to the files that are being closed. A more refined specification would be the following, requiring that a file can be closed only if it was opened before, and has not been closed since. Here, we use the temporal operators \ominus ("at previous step") and S ("since"):

$$\forall f(close(f) \longrightarrow \ominus(\neg close(f) \, \mathcal{S} \, open(f))) \tag{2}$$

One problem we need to solve is the unboundedness caused by negation. For example, assume that we have only observed so far one *close* event *close*("ab"). The subformula close(f) is therefore satisfied for the value f = "ab". The subformula $\neg close(f)$ is satisfied by all values from the domain of f except for "ab". This set contains those values that we have not seen yet in the input within a *close* event. We need a representation of finite and infinite sets of values, upon which applying complementation is efficient. We present a first-order past time temporal logic, named QTL (Quantified Temporal Logic), and an implementation, named DEJAVU based on a BDD (Binary Decision Diagram) representation of sets of assignments of values to the free variables of subformulas.

5.1 Syntax

Assume a finite set of domains D_1, D_2, \ldots Assume further that the domains are infinite, e.g., they can be the integers or strings⁷. Let *V* be a finite set of *variables*, with typical instances *x*, *y*, *z*. An *assignment* over a set of variables *W* maps each variable $x \in W$ to a value from its associated domain *domain*(*x*). For example $[x \rightarrow 5, y \rightarrow \text{``abc''}]$ maps *x* to 5 and *y* to "abc". Let *T* be a set of *predicate names* with typical instances *p*, *q*, *r*. Each predicate name *p* is associated with some domain *domain*(*p*). A predicate is constructed from a predicate name, and a variable or a constant of the same type. Thus, if the predicates like p(``gaga''), p(``baba'') and p(x). Predicates over constants are called *ground predicates*. An *event* is a finite set of ground predicates. For example, if $T = \{p, q, r\}$, then $\{p(\text{``xyzzy''}), q(3)\}$ is a possible event. An *execution* $\sigma = s_1s_2...$ is a finite sequence of events.

For runtime verification, a property φ is interpreted on prefixes of a monitored sequence. We check whether φ holds for every such prefix, hence, conceptually, check whether $\Box \varphi$ holds, where \Box is the "always in the future" linear temporal logic operator. The formulas of the core logic QTL are defined by the following grammar. For simplicity of the presentation, we define here the logic with unary predicates, but this is not due to any principle limitation, and, in fact, our implementation supports predicates with multiple arguments.

$$\varphi ::= true \mid p(a) \mid p(x) \mid (\varphi \land \varphi) \mid \neg \varphi \mid (\varphi \mathrel{\mathcal{S}} \varphi) \mid \ominus \varphi \mid \exists x \varphi$$

The formula p(a), where *a* is a constant in *domain*(*p*), means that the ground predicate p(a) occurs in the most recent event. The formula p(x), for a variable $x \in V$, holds with a binding of *x* to the value *a* if a ground predicate p(a) appears in the most recent event. The formula $(\varphi_1 \ S \ \varphi_2)$ means that φ_2 held in the past (possibly now) and since then φ_1 has been true. The property $\ominus \varphi$ means that φ was true in the previous event. We can also define the following additional operators: *false* = $\neg true$, ($\varphi \lor \psi$) = $\neg(\neg \varphi \land \neg \psi)$, ($\varphi \longrightarrow \psi$) = ($\neg \varphi \lor \psi$), **P** φ = (*true* $S \ \varphi$) (*previously* φ), **H** $\varphi = \neg \mathbf{P} \neg \varphi$ (*historically* φ , or φ *always in the past*), and $\forall x \ \varphi = \neg \exists x \neg \varphi$. The operator [φ, ψ), borrowed from [19], has the same meaning as ($\neg \psi \ S \ \varphi$), but reads more naturally as an interval.

5.2 Semantics

Predicate Semantics

Let $free(\varphi)$ be the set of free (i.e., unquantified) variables of a subformula φ . Then $(\gamma, \sigma, i) \models \varphi$, where γ is an assignment over $free(\varphi)$, and $i \ge 1$, if φ holds for the prefix $s_1s_2...s_i$ of the execution σ with the assignment γ . We denote by $\gamma|_{free(\varphi)}$ the restriction (projection) of an assignment γ to the free variables appearing in φ and by ε the empty assignment. The semantics of QTL can be defined as follows.

- $(\varepsilon, \sigma, i) \models true$.

⁷ For dealing with finite domains see [15].

- $(\varepsilon, \sigma, i) \models p(a)$ if $p(a) \in \sigma[i]$.
- $([x \mapsto a], \sigma, i) \models p(x)$ if $p(a) \in \sigma[i]$.
- $(\gamma, \sigma, i) \models (\phi \land \psi)$ if $(\gamma|_{free(\phi)}, \sigma, i) \models \phi$ and $(\gamma|_{free(\psi)}, \sigma, i) \models \psi$.
- $(\gamma, \sigma, i) \models \neg \phi$ if not $(\gamma, \sigma, i) \models \phi$.
- $(\gamma, \sigma, i) \models (\phi S \Psi)$ if $(\gamma|_{free(\Psi)}, \sigma, i) \models \Psi$ or the following hold: i > 1, $(\gamma|_{free(\Phi)}, \sigma, i) \models \phi$, and $(\gamma, \sigma, i 1) \models (\phi S \Psi)$.
- $(\gamma, \sigma, i) \models \ominus \varphi$ if i > 1 and $(\gamma, \sigma, i 1) \models \varphi$.
- $(\gamma, \sigma, i) \models \exists x \phi$ if there exists $a \in domain(x)$ such that⁸ $(\gamma[x \mapsto a], \sigma, i) \models \phi$.

Set Semantics

It helps to present the BDD-based algorithm by first redefining the semantics of the logic in terms of sets of assignments satisfying a formula. Let $I[\phi, \sigma, i]$ be the semantic function, defined below, that returns a set of assignments such that $\gamma \in I[\phi, \sigma, i]$ iff $(\gamma, \sigma, i) \models \phi$. The empty set of assignments \emptyset behaves as the Boolean constant 0 and the singleton set that contains an assignment over an empty set of variables $\{\varepsilon\}$ behaves as the Boolean constant 1. We define the union and intersection operators on sets of assignments, even if they are defined over non identical sets of variables. In this case, the assignments are extended over the union of the variables. Thus intersection between two sets of assignments A_1 and A_2 is defined like a database "join" operator; i.e., it consists of the assignments whose projection on the *common* variables agrees with an assignment in A_1 and with an assignment in A_2 . Union is defined as the dual operator of intersection. Let A be a set of assignments over the set of variables W; we denote by hide(A, x) (for "hiding" the variable x) the set of assignments obtained from A after removing from each assignment the mapping from x to a value. In particular, if A is a set of assignments over only the variable x, then hide(A, x) is $\{\varepsilon\}$ when A is nonempty, and \emptyset otherwise. $A_{free(\varphi)}$ is the set of all possible assignments of values to the variables that appear free in φ . We add a 0 position for each sequence σ (which starts with s_1), where *I* returns the empty set for each formula. The assignment-set semantics of QTL is shown in the following. For all occurrences of *i*, it is assumed that i > 0.

 $\begin{array}{l} - I[\varphi, \sigma, 0] = \emptyset. \\ - I[true, \sigma, i] = \{\varepsilon\}. \\ - I[p(a), \sigma, i] = \text{if } p(a) \in \sigma[i] \text{ then } \{\varepsilon\} \text{ else } \emptyset. \\ - I[p(x), \sigma, i] = \{[x \mapsto a] | p(a) \in \sigma[i]\}. \\ - I[(\varphi \land \psi), \sigma, i] = I[\varphi, \sigma, i] \cap I[\psi, \sigma, i]. \\ - I[\neg \varphi, \sigma, i] = A_{free(\varphi)} \setminus I[\varphi, \sigma, i]. \\ - I[(\varphi \mathrel{\mathcal{S}} \psi), \sigma, i] = I[\psi, \sigma, i] \cup (I[\varphi, \sigma, i] \cap I[(\varphi \mathrel{\mathcal{S}} \psi), \sigma, i-1]). \\ - I[\ominus \varphi, \sigma, i] = I[\varphi, \sigma, i-1]. \\ - I[\exists x \varphi, \sigma, i] = hide(I[\varphi, \sigma, i], x). \end{array}$

As before, the interpretation for the rest of the operators can be obtained from the above using the connections between the operators.

⁸ $\gamma[x \mapsto a]$ is the overriding of γ with the binding $[x \mapsto a]$.

6 Runtime Verification for First-Order Past LTL

We describe an algorithm for monitoring QTL properties, first presented in [15] and implemented in the tool DEJAVU. To give a brief overview of the contents of this section, instead of storing the data values occurring in events, we enumerate these data values as soon as we see them and use Boolean encodings of this enumeration. We use BDDs to represent sets of such enumerations. For example, if the runtime verifier sees the input events *open*("a"), *open*("b"), *open*("c"), it will encode them as 000, 001 and 010 (say, we use 3 bits b_0 , b_1 and b_2 to represent each enumeration, with b_2 being the most significant bit). A BDD that represents the set of values {"a", "c"} would be equivalent to a Boolean function ($\neg b_0 \land \neg b_2$) that returns 1 for 000 and 010 (the value of b_1 can be arbitrary). This approach has the following benefits:

- It is highly compact. With k bits used for representing enumerations, the BDD can grow to $2^{O(k)}$ nodes [8]; but BDDs usually compact the representation very well [10]. In fact, we often do not pay much in overhead for keeping surplus bits. Thus, we can start with an overestimated number of bits k such that it is unlikely to see more than 2^k different values for the domain they represent. We can also incrementally extend the BDD with additional bits when needed at runtime.
- Complementation (negation) is efficient, by just switching between the 0 and 1 leaves of the BDD. Moreover, even though at any point we may have not seen the entire set of values that will show up during the execution, we can safely (and efficiently) perform complementation: values that have not appeared yet in the execution are being accounted for and their enumerations are reserved already in the BDD before these values appear.
- Our representation of sets of assignments as BDDs allows a very simple algorithm that naturally extends the dynamic programming monitoring algorithm for propositional past time temporal logic shown in [17] and summarized in Setion 4.2.

6.1 BDDs

We represent a set of assignments as an Ordered Binary Decision Diagram (OBDD, although we write simply BDD) [9]. A BDD is a compact representation for a Boolean valued function of type $\mathbb{B}^k \to \mathbb{B}$ for some k > 0 (where \mathbb{B} is the Boolean domain $\{0, 1\}$), as a directed acyclic graph (DAG). A BDD is essentially a compact representation of a Boolean tree, where compaction glues together isomorphic subtrees. Each non-leaf node is labeled with one of the Boolean variables b_0, \ldots, b_{k-1} . A non-leaf node b_i is the source of two arrows leading to other nodes. A dotted-line arrow represents that b_i has the Boolean value 0, while a thick-line arrow represents that it has the value 1. The nodes in the DAG have the same order along all paths from the root. However, some of the nodes may be absent along some paths, when the result of the Boolean function does not depend on the value of the corresponding Boolean variable. Each path leads to a leaf node that is marked by either a 0 or a 1, representing the Boolean value returned by the function for the Boolean values on the path. Figure 3 contains five BDDs (a)-(e), over three Boolean variables b_0, b_1 , and b_2 (referred to by their subscripts 0, 1, and 2), as explained below.

6.2 Mapping data to BDDs

Assume that we see p("ab"), p("de"), p("af") and q("fg") in subsequent states in a trace, where p and q are predicates over the domain of strings. When a value associated with a variable appears for the first time in the current event (in a ground predicate), we add it to the set of values of that domain that were seen. We assign to each new value an *enumeration*, represented as a binary number, and use a hash table to point from the value to its enumeration.

Consistent with the DEJAVU implementation, the least significant bit in an enumeration is denoted in Figure 3 (and in the rest of this paper) by BDD variable 0, and the most significant bit by BDD variable n - 1, where n is the number of bits. Using e.g. a three-bit enumeration $b_2b_1b_0$, the first encountered value "ab" can be represented as the bit string 000, "de" as 001, "af" as 010 and "fg" as 011. A BDD for a subset of these values returns a 1 for each bit string representing an enumeration of a value in the set, and 0 otherwise. E.g. a BDD representing the set {"de", "af"} (2nd and 3rd values) returns 1 for 001 and 010. This is the Boolean function $\neg b_2 \land (b_1 \leftrightarrow \neg b_0)$. Figure 3 shows the BDDs for each of these values as well as the BDD for the set containing the values "de" and "af".

When representing a set of assignments for e.g. two variables *x* and *y* with *k* bits each, we will have Boolean variables $x_0, \ldots, x_{k-1}, y_0, \ldots, y_{k-1}$. A BDD will return a 1 for each bit string representing the concatenation of enumerations that correspond to the represented assignments, and 0 otherwise. For example, to represent the assignments $[x \mapsto "de", y \mapsto "af"]$, where "de" is enumerated as 001 and "af" with 010, the BDD will return a 1 for 001010.

6.3 The BDD-based algorithm

Given some ground predicate p(a) observed in the execution matching with p(x) in the monitored property, let **lookup**(x, a) be the enumeration of a. If this is a's first occurrence, then it will be assigned a new enumeration. Otherwise, **lookup** returns the enumeration that a received before. We can use a counter, for each variable x, counting the number of different values appearing so far for x. When a new value appears, this counter is incremented, and the value is converted to a Boolean representation. Enumerations that were not yet used represent the values not seen yet. In the next section we introduce data reclaiming, which allows reusing enumerations for values that no longer affect the checked property. This involves a more complicated enumeration mechanism.

The function **build**(*x*,*A*) returns a BDD that represents the set of assignments where *x* is mapped to (the enumeration of) *v* for $v \in A$. This BDD is independent of the values assigned to any variable other than *x*, i.e., they can have any value. For example, assume that we use three Boolean variables (bits) x_0 , x_1 and x_2 for representing enumerations over *x* (with x_0 being the least significant bit), and assume that $A = \{a, b\}$, **lookup**(x, a) = 011, and **lookup**(x, b) = 001. Then **build**(x, A) is a BDD representation of the Boolean function $x_0 \land \neg x_2$.

Intersection and union of sets of assignments are translated simply to conjunction and disjunction of their BDD representation, respectively, and complementation becomes BDD negation. We will denote the Boolean BDD operators as **and**, **or** and **not**.



Fig. 3: BDDs for the trace: *p*("ab").*p*("de").*p*("af").*q*("fg")

To implement the existential (universal, respectively) operators, we use the BDD existential (universal, respectively) operators over the Boolean variables that represent (the enumerations of) the values of *x*. Thus, if B_{φ} is the BDD representing the assignments satisfying φ in the current state of the monitor, then **exists**($\langle x_0, \ldots, x_{k-1} \rangle, B_{\varphi}$) is the BDD that represents the assignments satisfying $\exists x \varphi$ in the current state. Finally, BDD(\perp) and BDD(\top) are the BDDs that return always 0 or 1, respectively.

The dynamic programming algorithm, shown below, works similarly to the algorithm for the propositional case shown in Section 4.2. That is, it operates on two vectors (arrays) of values indexed by subformulas: pre for the state before that event, and now for the current state (after the last seen event). However, while in the propositional case the vectors contain Boolean values, here they contain BDDs. The algorithm follows.

- 1. Initially, for each subformula φ , now(φ) := BDD(\perp).
- 2. Observe a new event (as a set of ground predicates) s as input.
- 3. Let pre := now.
- 4. Make the following updates for each subformula. If ϕ is a subformula of ψ then $\mathsf{now}(\phi)$ is updated before $\mathsf{now}(\psi)$.
 - now(*true*) := BDD(\top).
 - now $(p(a)) := \text{if } p(a) \in s \text{ then } BDD(\top) \text{ else } BDD(\bot).$
 - now(p(x)) := **build**(x, A) where $A = \{a | p(a) \in s\}$.
 - $\operatorname{now}((\phi \land \psi)) := \operatorname{and}(\operatorname{now}(\phi), \operatorname{now}(\psi)).$
 - $now(\neg \phi) := not(now(\phi)).$
 - now(($\varphi \mathcal{S} \psi$)) := or(now(ψ), and(now(φ), pre(($\varphi \mathcal{S} \psi$)))).
 - now($\ominus \phi$) := pre(ϕ).
 - now($\exists x \varphi$) := exists($\langle x_0, \ldots, x_{k-1} \rangle$, now(φ)).
- 5. Goto step 2.

An important property of the algorithm is that, at any point during monitoring, enumerations that are not used in the pre and now BDDs represent all values that have *not* been seen so far in the input. This can be proved by induction on the size of temporal formulas and the length of the input sequence. We specifically identify one enumeration to represent all values not seen yet, namely the largest possible enumeration, given the number of bits we use, 11...11. We let BDD(11...11) denote the BDD that returns 1 exactly for this value. This trick allows us to use a finite representation and quantify existentially and universally over *all* values in infinite domains.

6.4 An Example

DEJAVU is implemented in SCALA. DEJAVU takes as input a specification file containing one or more properties, and synthesizes a self-contained SCALA program (a text file) - the monitor, as already illustrated for the propositional case in Section 4.2. This program (which first must be compiled) takes as input the trace file and analyzes it. The tool uses the JavaBDD library for BDD manipulations [18]. We shall illustrate the monitor generation using an example. Consider the following property stating that if a file f is closed, it must have been opened in the past with some access mode m (e.g. 'read' or 'write' mode):

$$\forall f(close(f) \longrightarrow \exists m \mathbf{P} open(f,m))$$

The property-specific part⁹ of the synthesized monitor, shown in Figure 4 (left), relies on the enumeration of the subformulas shown in Figure 4 (right). As in the propositional case, two arrays are declared, indexed by subformula indexes: pre for the previous state and now for the current state, although here storing BDDs instead of Boolean values. For each observed event, the function evaluate() computes the now array from highest to lowest index, and returns true (property is satisfied in this position of the trace) iff now(0) is not BDD(\perp). At composite subformula nodes, BDD operators are applied.

⁹ An additional 600+ lines of property independent boilerplate code is generated.

For example for subformula 4, the new value is now(5).or(pre(4)), which is the interpretation of the formula **P** open(f, m) corresponding to the law: $\mathbf{P}\phi = (\phi \lor \ominus \mathbf{P}\phi)$. As can be seen, for each new event, the evaluation of a formula results in the computation of a BDD for each subformula.



Fig. 4: Monitor (left) and AST (right) for the property

We shall briefly evaluate the example formula on a trace. Assume that each variable f and m is represented by three BDD bits. Consider the input trace, consisting of three events¹⁰:

open(input,read).open(output,write).close(out)

When the monitor evaluates subformula 5 on the first event *open*(input, read), it will create a bit string composed of a bit string for each parameter f and m. As previously explained, bit strings for each variable are allocated in increasing order: 000, 001, 010,..., hence the first bit string representing the assignment [$f \mapsto$ input, $m \mapsto$ read] becomes 000000 where the three rightmost bits represent the assignment of input to f, and the three leftmost bits represent the assignment of read to m. Figure 5a shows the corresponding BDD B_1 . Recall that most significant bits are implemented lower in the BDD, and that for each bit (node) in the BDD, the dotted arrow corresponds to this bit being 0 and the full drawn arrow corresponds to this bit being 1. In this BDD all bits have to be zero in order to be accepted by the function represented by the BDD. We will not show

¹⁰ Traces accepted by the tool are concretely in CSV format. For example the first event is a single line of the form: open, input, read.



Fig. 5: Selected BDDs, named B_1, \ldots, B_6 , computed after each event at various subformula nodes, indicated by BDD B_i @ node (see Figure 4), during processing of the trace: *open*(input,read).*open*(output,write).*close*(out).

how all the tree nodes evaluate, except observe that node 5 assumes the same BDD value as node 4 (all the seen values in the past), and conclude that since no close(...) event has been observed, the top-level formula (node 0) is true at this position in the trace.

Upon the second *open*(output,write) event, new values (output,write) are observed as argument to the *open* event. Hence a new bit string for each variable f and m is allocated, in both cases 001 (the next unused bit string for each variable). The new combined bit string for the assignments satisfying subformula 5 then becomes 001001, forming a BDD representing the assignment [$f \mapsto output$, $m \mapsto write$], and appearing in Figure 5b as B_2 . The computation of the BDD for node 4 is computed by now(4) = now(5).or(pre(4)), which results in the BDD B_3 , representing the set of the two so far observed assignments ($B_3 = or(B_1, B_2)$).

Upon the third *close*(out) event, a new value out for f is observed, and allocated the bit pattern 010, represented by the BDD B_4 for subformula 2. At this point node 4 still evaluates to the BDD B_3 (unchanged from the previous step), and the existential quantification over m in node 3 results in the BDD B_5 , where the bits 3, 4 and 5 for m have been removed, and the BDD compacted. Node 1 is computed as **or**(**not**(B_4), B_5), which results in the BDD B_6 . This BDD represents all bit patterns for f that are **not** 010, corresponding to the value: out. So for all such values the formula is true. This means, however, that the top-level formula in node 0 is not true (violated by bit pattern 010), and hence the formula is violated on the third event.

6.5 Dynamic Data Reclamation

Consider property 1 on page 10 that asserts that each file that is closed was opened before. If we do not remember for this property *all* the files that were opened, then we will not be able to check when a file is closed whether it was opened before. Consider now the more refined property 2 on page 10, requiring that a file can be closed only if it was opened before, and has not been closed. Observe here that if a file was opened and subsequently closed, then if it is closed again before opening, the property would be invalidated just as in the case where it was not opened at all. This means that we can "forget" that a file was opened when it is closed without affecting our ability to monitor the formula. This allows reusing enumerations of data values, when this does not affect the decision whether the property holds or not.

Let *A* be a set of assignments over some variables that include *x*. Denote by A[x = a] the set of assignments from *A* in which the value of *x* is *a*. We say that the values *a* and *b* are *analogous* for variable *x* in *A*, if *hide*(A[x = a], x) = *hide*(A[x = b], x). This means that *a* and *b*, as values of the variable *x*, are related to all other values in *A* in the same way. A value can be reclaimed if it is analogous to the values not seen yet in all the assignments represented in pre(ψ), for each subformula ψ .

We shall now identify enumerations that can be reclaimed, and remove the values in the hash table that map to them, such that the enumerations can later be reused to represent new values. The search for reclaimable enumerations in a particular step during monitoring is performed on the pre BDDs. Recall that the enumeration 11...11 represents all the values that were *not* seen so far. Thus, we can check whether a value *a* for *x* is analogous to the values not seen so far for *x* by performing the checks on the pre BDDs between the enumeration of *a* and the enumeration 11...11. In fact, we do not have to perform the checks enumeration by enumeration, but use a BDD expression that constructs a BDD representing (returning 1 for) all enumerations that can be reclaimed for a variable *x*.

Assume that a subformula ψ has three free variables, x, y and z, each with k bits, i.e., $x_0, \ldots, x_{k-1}, y_0, \ldots, y_{k-1}$ and z_0, \ldots, z_{k-1} . The following expression returns a BDD representing the enumerations for values of x in assignments represented by pre(ψ) that are related to enumerations of y and z in the same way as 11...11.

$$I_{\Psi,x} = \forall y_0 \dots \forall y_{k-1} \forall z_0 \dots \forall z_{k-1} (\mathsf{pre}(\Psi)[x_0 \setminus 1, \dots, x_{k-1} \setminus 1] \leftrightarrow \mathsf{pre}(\Psi))$$

We now conjoin the above formula over each subformula that has a temporal operator at the outermost level, and subtract from this conjunction the 11...11 enumeration. This becomes the BDD *avail* of available enumerations. Any enumeration that is in *avail* can be reclaimed, and later reused as the enumeration of a new value. The selection of a "free" enumeration from *avail* can be performed by a SAT solver that picks any enumeration that is in *olonger* available. Note that if a value later reappears after its enumeration was reclaimed, it is likely to be assigned a different enumeration.

7 Conclusion

We presented a collection of runtime verification algorithms for linear temporal logics. First we introduced propositional temporal logic, including future as well as past time operators. We presented a theory of monitorability of temporal properties, introducing classes that reflect different degrees of monitorability. The notion of monitorability identifies the kinds of verdicts that one can obtain from observing finite prefixes of an execution. We then presented monitoring algorithms for the future time case as automata, and for the past time case as an instance of dynamic programming. We also provided algorithms for checking what kind of verdict (positive or negative) we can expect after monitoring a certain prefix against a given property, and whether a property is monitorable or not. We then introduced first-order past time linear temporal logic, and a monitoring algorithm for it. While the propositional version is independent of the length of the prefix seen so far, the first-order version may need to represent an amount of values that can grow linearly with the number of data values observed so far. The challenge is to provide a compact representation that will grow slowly and can be updated quickly with each incremental calculation that is performed per each new monitored event. We used a BDD representation of sets of assignments for the variables that appear in the monitored property.

References

- B. Alpern, F. B. Schneider, Recognizing Safety and Liveness. Distributed Computing 2(3), 117-126, 1987.
- E. Bartocci, Y. Falcone, A. Francalanza, M. Leucker, G. Reger, An Introduction to Runtime Verification, Lectures on Runtime Verification - Introductory and Advanced Topics, LNCS Volume 10457, Springer, 1-23, 2018.
- D. A. Basin, C. C. Jiménez, F. Klaedtke, E. Zalinescu, Deciding Safety and Liveness in TPTL. Information Processing Letters 114(12), 680-688, 2014.
- D. A. Basin, F. Klaedtke, S. Müller, E. Zalinescu, Monitoring Metric First-Order Temporal Properties, Journal of the ACM 62(2), 1-45, 2015.
- A. Bauer, J.C. Küster, G. Vegliach, From Propositional to First-Order Monitoring, RV'13, LNCS Volume 8174, Springer, 59-75, 2013.
- A. Bauer, M. Leucker, C. Schallhart, The Good, the Bad, and the Ugly, But How Ugly is Ugly?, RV'07, LNCS Volume 4839, Springer, 126-138, 2007.

- A. Bauer, M. Leucker, C. Schallhart, Runtime verification for LTL and TLTL. ACM Trans. Software Engineering Methodologies, 20(4), 14:1-14:64, 2011.
- R. E. Bryant, On the Complexity of VLSI Implementations and Graph Representations of Boolean Functions with Application to Integer Multiplication, IEEE Transactions on Computers 40(2): 205-213 (1991).
- R. E. Bryant, Symbolic Boolean Manipulation with Ordered Binary-Decision Diagrams, ACM Comput. Surv. 24(3), 293-318, 1992.
- J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, L. J. Hwang, Symbolic Model Checking: 10²⁰ States and Beyond, LICS'90, 428-439, 1990.
- Y. Falcone, J.-C. Fernandez, L. Mounier, Runtime Verification of Safety/Progress Properties, RV'09, LNCS Volume 5779, Springer, 40-59, 2009.
- Y. Falcone, J.-C. Fernandez, L. Mounier, What can you Verify and Enforce at Runtime?, STTT 14(3), 349-382, 2012.
- D. Gabbay, A. Pnueli, S. Shelah, J. Stavi, On the Temporal Analysis of Fairness, POPL'80, ACM, 163-173 1980.
- R. Gerth, D. A. Peled, M. Y. Vardi, P. Wolper, Simple On-The-Fly Automatic Verification of Linear Temporal Logic, PSTV'95, 3-18, 1995.
- K. Havelund, D. Peled, D. Ulus, First-order Temporal Logic Monitoring with BDDs, FM-CAD'17, IEEE, 116-123, 2017.
- K. Havelund, G. Reger, D. Thoma, E. Zălinescu, Monitoring Events that Carry Data, book chapter in: Lectures on Runtime Verification - Introductory and Advanced Topics, book editors: Ezio Bartocci and Yliès Falcone, LNCS Volume 10457, Springer, 61-102, 2018.
- K. Havelund, G. Rosu, Synthesizing Monitors for Safety Properties, TACAS'02, LNCS Volume 2280, Springer, 342-356, 2002.
- 18. JavaBDD, http://javabdd.sourceforge.net.
- M. Kim, S. Kannan, I. Lee, O. Sokolsky, Java-MaC: a Run-time Assurance Tool for Java, RV'01, Elsevier, ENTCS 55(2), 218-235, 2001.
- O. Kupferman, M. Y. Vardi, Model Checking of Safety Properties. Formal Methods in System Design 19(3): 291-314, 2001.
- O. Kupferman, G. Vardi, On Relative and Probabilistic Finite counterability. Formal Methods in System Design 52(2): 117-146, 2018.
- L. Lamport, Proving the Correctness of Multiprocess Programs, IEEE Transactions on Software Engineering 3(2): 125-143, 1977.
- Z. Manna, A. Pnueli, Completing the Temporal Picture, Theoretical Computer Science 83, 91-130, 1991.
- N. Markey, Temporal Logic with Past is Exponentially More Succinct, Concurrency Column. Bulletin of the EATCS 79: 122-128 (2003).
- P. O. Meredith, D. Jin, D. Griffith, F. Chen, G. Rosu, An Overview of the MOP Runtime Verification Framework, STTT 14(3), Springer, 249-289, 2012.
- D. Peled, K. Havelund, Refining the Safety–liveness Classification of Temporal Properties According to Monitorability, Submitted for publication, LNCS, Sept. 2018.
- A. Pnueli, A. Zaks, PSL Model Checking and Run-Time Verification via Testers. FM'06, LNCS Volume 4085, Springer, 573-586, 2006.
- A. P. Sistla, Safety, Liveness and Fairness in Temporal Logic, Formal Aspects of Computing 6(5): 495-512, 1994.
- A. P. Sistla, E. M. Clarke, The Complexity of Propositional Linear Temporal Logics, STOC'82, 159-168, 1982. Journal of the ACM (JACM) JACM Homepage archive Volume 32 Issue 3, July 1985 Pages 733-749
- W. Thomas, Automata on Infinite Objects, Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics, 133-192, 1990.