

Algorithms for Monitoring Real-time Properties^{*}

David Basin, Felix Klaedtke, and Eugen Zălinescu

Computer Science Department, ETH Zurich, Switzerland

Abstract. We present and analyze monitoring algorithms for a safety fragment of metric temporal logics, which differ in their underlying time model. The time models considered have either dense or discrete time domains and are point-based or interval-based. Our analysis reveals differences and similarities between the time models for monitoring and highlights key concepts underlying our and prior monitoring algorithms.

1 Introduction

Real-time logics [2] allow us to specify system properties involving timing constraints, e.g., every request must be followed within 10 seconds by a grant. Such specifications are useful when designing, developing, and verifying systems with hard real-time requirements. They also have applications in runtime verification, where monitors generated from specifications are used to check the correctness of system behavior at runtime [10].

Various monitoring algorithms for real-time logics have been developed [4, 5, 7, 12, 14, 15, 17, 20] based on different time models. These time models can be characterized by two independent aspects. First, a time model is either point-based or interval-based. In point-based time models, system traces are sequences of system states, where each state is time-stamped. In interval-based time models, system traces consist of continuous (Boolean) signals of state variables. Second, a time model is either dense or discrete depending on the underlying ordering on time-points, i.e., whether there are infinitely many or finitely many time-points between any two distinct time-points.

Real-time logics based on a dense, interval-based time model are more natural and general than their counterparts based on a discrete or point-based model. In fact, both discrete and point-based time models can be seen as abstractions of dense, interval-based time models [2, 18]. However, the satisfiability and the model-checking problems for many real-time logics with the more natural time model are computationally harder than their corresponding decision problems when the time model is discrete or point-based. See the survey [16] for further discussion and examples.

In this paper, we analyze the impact of different time models on monitoring. We do this by presenting, analyzing, and comparing monitoring algorithms for real-time logics based on different time models. More concretely, we present

^{*} This work was supported by the Nokia Research Center, Switzerland.

monitoring algorithms for the past-only fragment of propositional metric temporal logics with a point-based and an interval-based semantics, also considering both dense and discrete time domains. We compare our algorithms on a class of formulas for which the point-based and the interval-based settings coincide. To define this class, we distinguish between event propositions and state propositions. The truth value of a state proposition always has a duration, whereas an event proposition cannot be continuously true between two distinct time-points.

Our analysis explains the impact of different time models on monitoring. First, the impact of a dense versus a discrete time domain is minor. The algorithms are essentially the same and have almost identical computational complexities. Second, monitoring in a point-based setting is simpler than in an interval-based setting. The meaning of “simpler” is admittedly informal here since we do not provide lower bounds. However, we consider our monitoring algorithms for the point-based setting as conceptually simpler than the interval-based algorithms. Moreover, we show that our point-based monitoring algorithms perform better than our interval-based algorithms on the given class of formulas on which the two settings coincide.

Overall, we see the contributions as follows. First, our monitoring algorithms simplify and clarify key concepts of previously presented algorithms [4, 13–15]. In particular, we present the complete algorithms along with a detailed complexity analysis for monitoring properties specified in the past-only fragment of propositional metric temporal logic. Second, our monitoring algorithm for the dense, point-based time model has better complexity bounds than existing algorithms for the same time model [20]. Third, our comparison of the monitoring algorithms illustrates the similarities, differences, and trade-offs between the time models with respect to monitoring. Moreover, formulas in our fragment benefit from both settings: although they describe properties based on a more natural time model, they can be monitored with respect to a point-based time model, which is more efficient.

The remainder of the paper is structured as follows. In Section 2, we give preliminaries. In Section 3, we compare the point-based and the interval-based time model and define our class of formulas on which the two time models coincide. In Section 4, we present, analyze and compare our monitoring algorithms. In Section 5, we discuss related work. Finally, in Section 6, we draw conclusions. The appendices contain additional details.

2 Preliminaries

In this section, we fix the notation and terminology that we use in the remainder of the text.

Time Domain and Intervals. If not stated differently, we assume the dense time domain¹ $\mathbb{T} = \mathbb{Q}_{\geq 0}$ with the standard ordering \leq . Adapting the following definitions to a discrete time domain like \mathbb{N} is straightforward.

¹ We do not use $\mathbb{R}_{\geq 0}$ as dense time domain because of representation issues. Namely, each element in $\mathbb{Q}_{\geq 0}$ can be finitely represented, which is not the case for $\mathbb{R}_{\geq 0}$.

A *(time) interval* is a non-empty set $I \subseteq \mathbb{T}$ such that if $\tau < \kappa < \tau'$ then $\kappa \in I$, for all $\tau, \tau' \in I$ and $\kappa \in \mathbb{T}$. We denote the set of all time intervals by \mathbb{I} . An interval is either left-open or left-closed and similarly either right-open or right-closed. We denote the left margin and the right margin of an interval $I \in \mathbb{I}$ by $\ell(I)$ and $r(I)$, respectively. For instance, the interval $I = \{\tau \in \mathbb{T} \mid 3 \leq \tau\}$, which we also write as $[3, \infty)$, is left-closed and right-open with margins $\ell(I) = 3$ and $r(I) = \infty$.

For an interval $I \in \mathbb{I}$, we define the extension $I^{\geq} := I \cup (\ell(I), \infty)$ to the right and its strict counterpart $I^{>} := I^{\geq} \setminus I$, which excludes I . We define $I^{\leq} := [0, r(I)) \cup I$ and $I^{<} := (I^{\leq}) \setminus I$ similarly. An interval $I \in \mathbb{I}$ is *singular* if $|I| = 1$, *bounded* if $r(I) < \infty$, and *unbounded* if $r(I) = \infty$. The intervals $I, J \in \mathbb{I}$ are *adjacent* if $I \cap J = \emptyset$ and $I \cup J \in \mathbb{I}$. For $I, J \in \mathbb{I}$, $I \oplus J$ is the set $\{\tau + \tau' \mid \tau \in I \text{ and } \tau' \in J\}$.

An *interval partition* of \mathbb{T} is a sequence $\langle I_i \rangle_{i \in \mathbb{N}}$ of time intervals with $N = \mathbb{N}$ or $N = \{0, \dots, n\}$ for some $n \in \mathbb{N}$ that fulfills the following properties: (i) I_{i-1} and I_i are adjacent and $\ell(I_{i-1}) \leq \ell(I_i)$, for all $i \in N \setminus \{0\}$, and (ii) for each $\tau \in \mathbb{T}$, there is an $i \in N$ such that $\tau \in I_i$. The interval partition $\langle J_j \rangle_{j \in M}$ *refines* the interval partition $\langle I_i \rangle_{i \in N}$ if for every $j \in M$, there is some $i \in N$ such that $J_j \subseteq I_i$. We often write \bar{I} for a sequence of intervals instead of $\langle I_i \rangle_{i \in N}$. Moreover, we abuse notation by writing $I \in \langle I_i \rangle_{i \in N}$ if $I = I_i$, for some $i \in N$.

A *time sequence* $\langle \tau_i \rangle_{i \in \mathbb{N}}$ is a sequence of elements $\tau_i \in \mathbb{T}$ that is strictly increasing (i.e., $\tau_i < \tau_j$, for all $i, j \in \mathbb{N}$ with $i < j$) and progressing (i.e., for all $\tau \in \mathbb{T}$, there is $i \in \mathbb{N}$ with $\tau_i > \tau$). Similar to interval sequences, $\bar{\tau}$ abbreviates $\langle \tau_i \rangle_{i \in \mathbb{N}}$.

Boolean Signals. A *(Boolean) signal* γ is a subset of \mathbb{T} that fulfills the following finite-variability condition: for every bounded interval $I \in \mathbb{I}$, there are intervals $I_0, \dots, I_{n-1} \in \mathbb{I}$ such that $\gamma \cap I = I_0 \cup \dots \cup I_{n-1}$, for some $n \in \mathbb{N}$. The least such $n \in \mathbb{N}$ is the *size* of the signal γ on I . We denote it by $\|\gamma \cap I\|$.

We use the term “signal” for such a set γ because its characteristic function $\chi_\gamma : \mathbb{T} \rightarrow \{0, 1\}$ represents, for example, the values over time of an input or an output of a sequential circuit. Intuitively, $\tau \in \gamma$ iff the signal of the circuit is high at the time $\tau \in \mathbb{T}$. The finite-variability condition imposed on the set γ prevents switching infinitely often from high to low in finite time. Note that $\|\gamma \cap I\|$ formalizes how often a signal γ is high on the bounded interval I , in particular, $\|\gamma \cap I\| = 0$ iff $\gamma \cap I = \emptyset$.

A signal γ is *stable* on an interval $I \in \mathbb{I}$ if $I \subseteq \gamma$ or $I \cap \gamma = \emptyset$. The *induced interval partition* $\overline{\text{ip}}(\gamma)$ of a signal γ is the interval partition \bar{I} such that γ is stable on each of the intervals in \bar{I} and any other stable interval partition refines \bar{I} . We write $\overline{\text{ip}}^1(\gamma)$ for the sequence of intervals I in $\overline{\text{ip}}(\gamma)$ such that $I \cap \gamma \neq \emptyset$. Similarly, we write $\overline{\text{ip}}^0(\gamma)$ for the sequence of intervals I in $\overline{\text{ip}}(\gamma)$ such that $I \cap \gamma = \emptyset$. Intuitively, $\overline{\text{ip}}^1(\gamma)$ and $\overline{\text{ip}}^0(\gamma)$ are the sequences of maximal intervals on which the signal is γ is high and low, respectively.

Choosing $\mathbb{Q}_{\geq 0}$ instead of $\mathbb{R}_{\geq 0}$ is without loss of generality for the satisfiability of properties specified in real-time logics like the metric interval temporal logic [1].

$\begin{aligned} \hat{\gamma}, \tau \models p & \quad \text{iff } \tau \in \gamma_p \\ \hat{\gamma}, \tau \models \neg\phi & \quad \text{iff } \hat{\gamma}, \tau \not\models \phi \\ \hat{\gamma}, \tau \models \phi \wedge \psi & \quad \text{iff } \hat{\gamma}, \tau \models \phi \text{ and } \hat{\gamma}, \tau \models \psi \\ \hat{\gamma}, \tau \models \phi S_I \psi & \quad \text{iff there is } \tau' \in [0, \tau] \text{ with} \\ & \quad \tau - \tau' \in I, \\ & \quad \hat{\gamma}, \tau' \models \psi, \text{ and} \\ & \quad \hat{\gamma}, \kappa \models \phi, \text{ for all } \kappa \in (\tau', \tau] \end{aligned}$	$\begin{aligned} \hat{\gamma}, \bar{\tau}, i \models\dot{=} p & \quad \text{iff } \tau_i \in \gamma_p \\ \hat{\gamma}, \bar{\tau}, i \models\dot{=} \neg\phi & \quad \text{iff } \hat{\gamma}, \bar{\tau}, i \not\models\dot{=} \phi \\ \hat{\gamma}, \bar{\tau}, i \models\dot{=} \phi \wedge \psi & \quad \text{iff } \hat{\gamma}, \bar{\tau}, i \models\dot{=} \phi \text{ and } \hat{\gamma}, \bar{\tau}, i \models\dot{=} \psi \\ \hat{\gamma}, \bar{\tau}, i \models\dot{=} \phi S_I \psi & \quad \text{iff there is } i' \in [0, i] \cap \mathbb{N} \text{ with} \\ & \quad \tau_i - \tau_{i'} \in I, \\ & \quad \hat{\gamma}, \bar{\tau}, i' \models\dot{=} \psi, \text{ and} \\ & \quad \hat{\gamma}, \bar{\tau}, k \models\dot{=} \phi, \text{ for all } k \in (i', i] \cap \mathbb{N} \end{aligned}$
(a) interval-based semantics	(b) point-based semantics

Fig. 1. Semantics of past-only metric temporal logic.

Metric Temporal Logics. To simplify the exposition, we restrict ourselves to monitoring the past-only fragment of metric temporal logic in a point-based and an interval-based setting. However, future operators like \diamond_I , where the interval I is bounded, can be handled during monitoring by using queues that postpone the evaluation until enough time has elapsed. See [4], for such a monitoring algorithm that handles arbitrary nesting of past and bounded future operators.

Let P be a non-empty set of *propositions*. The syntax of the past-only fragment of metric temporal logic is given by the grammar $\phi ::= p \mid \neg\phi \mid \phi \wedge \phi \mid \phi S_I \phi$, where $p \in P$ and $I \in \mathbb{I}$. In Figure 1, we define the satisfaction relations \models and $\models\dot{=}$, where $\hat{\gamma} = (\gamma_p)_{p \in P}$ is a family of signals, $\bar{\tau}$ a time sequence, $\tau \in \mathbb{T}$, and $i \in \mathbb{N}$. Note that \models defines the truth value of a formula for every $\tau \in \mathbb{T}$. In contrast, a formula’s truth value with respect to $\models\dot{=}$ is defined at the “sample-points” $i \in \mathbb{N}$ to which the “time-stamps” $\tau_i \in \mathbb{T}$ from the time sequence $\bar{\tau}$ are attached.

We use the standard binding strength of the operators and standard syntactic sugar. For instance, $\phi \vee \psi$ stands for the formula $\neg(\neg\phi \wedge \neg\psi)$ and $\blacklozenge_I \psi$ stands for $(p \vee \neg p) S_I \psi$, for some $p \in P$. Moreover, we often omit the interval $I = [0, \infty)$ attached to a temporal operator. We denote the set of subformulas of a formula ϕ by $\text{sf}(\phi)$. Finally, $|\phi|$ is the number of nodes in ϕ ’s parse tree.

3 Point-based versus Interval-based Time Models

We first point out some shortcomings of a point-based time model in Section 3.1. In Section 3.2, we then present a class of formulas in which the point-based and the interval-based time models coincide.

3.1 State Variables and System Events

State variables and system events are different kinds of entities. One distinguishing feature is that events happen at single points in time and the value of a state variable is always constant for some amount of time. In the following, we distinguish between these two entities. Let P be the disjoint union of the proposition sets S and E . We call propositions in S *state propositions* and propositions in E *event propositions*. Semantically, a signal $\gamma \subseteq \mathbb{T}$ is an *event signal* if $\gamma \cap I$ is finite, for every bounded interval I , and the signal γ is a *state signal* if for every bounded interval I , the sets $\gamma \cap I$ and $(\mathbb{T} \setminus \gamma) \cap I$ are the finite unions of non-singular intervals. Note that there are signals that are neither event signals nor

state signals. A family of signals $\hat{\gamma} = (\gamma_p)_{p \in S \cup E}$ is *consistent* with S and E if γ_p is a state signal, for all $p \in S$, and γ_p is an event signal, for all $p \in E$.

The point-based semantics is often motivated by the study of real-time systems whose behavior is determined by system events. Intuitively, a time sequence $\bar{\tau}$ records the points in time when events occur and the signal γ_p for a proposition $p \in E$ consists of the points in time when the event p occurs. The following examples, however, demonstrate that the point-based semantics can be unintuitive in contrast to the interval-based semantics.

Example 1. A state proposition $p \in S$ can often be mimicked by the formula $\neg f \text{ S } s$ with corresponding event propositions $s, f \in E$ representing “start” and “finish.” For the state signal γ_p , let γ_s and γ_f be the event signals where γ_s and γ_f consist of the points in time of γ_p when the Boolean state variable starts and respectively finishes to hold. Then $(\gamma_s, \gamma_f), \tau \models \neg f \text{ S } s$ iff $\gamma_p, \tau \models p$, for any $\tau \in \mathbb{T}$, under the assumption that $I \cap \gamma_p$ is the finite union of left-closed and right-open intervals, for every bounded left-closed and right-open interval I .

However, replacing p by $\neg f \text{ S } s$ does not always capture the essence of a Boolean state variable when using the point-based semantics. Consider the formula $\blacklozenge_{[0,1]} p$ containing the state proposition p and let $\gamma_p = [0, 5)$ be a state signal. Moreover, let (γ_s, γ_f) be the family of corresponding event signals for the event propositions s and f , i.e., $\gamma_s = \{0\}$ and $\gamma_f = \{5\}$. For a time sequence $\bar{\tau}$ with $\tau_0 = 0$ and $\tau_1 = 5$, we have that $(\gamma_s, \gamma_f), \bar{\tau}, 1 \not\models \blacklozenge_{[0,1]} (\neg f \text{ S } s)$ but $\gamma_p, \tau_1 \models \blacklozenge_{[0,1]} p$. Note that $\bar{\tau}$ only contains time-stamps when an event occurs. An additional sample-point between τ_0 and τ_1 with, e.g., the time-stamp 4 would result in identical truth values at time 5.

Even when restricted to events, the point-based semantics can be unintuitive.

Example 2. Consider the (event) signals $\gamma_p = \{\tau \in \mathbb{T} \mid \tau = 2n, \text{ for some } n \in \mathbb{N}\}$ and $\gamma_q = \emptyset$ for the (event) propositions p and q . One might expect that these signals satisfy the formula $p \rightarrow \blacklozenge_{[0,1]} \neg q$ at every point in time. However, for a time sequence $\bar{\tau}$ with $\tau_0 = 0$ and $\tau_1 = 2$, we have that $\hat{\gamma}, \bar{\tau}, 1 \not\models p \rightarrow \blacklozenge_{[0,1]} \neg q$. The reason is that in the point-based semantics, the \blacklozenge_I operator requires the existence of a previous point in time that also occurs in the time sequence $\bar{\tau}$.

As another example consider the formula $\blacklozenge_{[0,1]} \blacklozenge_{[0,1]} p$. One might expect that it is logically equivalent to $\blacklozenge_{[0,2]} p$. However, this is not the case in the point-based semantics. To see this, consider a time sequence $\bar{\tau}$ with $\tau_0 = 0$ and $\tau_1 = 2$. We have that $\hat{\gamma}, \bar{\tau}, 1 \not\models \blacklozenge_{[0,1]} \blacklozenge_{[0,1]} p$ and $\hat{\gamma}, \bar{\tau}, 1 \models \blacklozenge_{[0,2]} p$ if $\tau_0 \in \gamma_p$.

The examples above suggest that adding additional sample-points restores a formula’s intended meaning, which usually stems from having the interval-based semantics in mind. However, a drawback of this approach for monitoring is that each additional sample-point increases the workload of a point-based monitoring algorithm, since it is invoked for each sample-point. Moreover, in the dense time domain, adding sample-points does not always make the two semantics coincide. For instance, for $\gamma_p = [0, 1)$ and $\tau \geq 1$, we have that $\gamma_p, \tau \not\models \neg p \text{ S } p$ and $\gamma_p, \bar{\tau}, i \models \neg p \text{ S } p$, for every time sequence $\bar{\tau}$ with $\tau_0 < 1$ and every $i \in \mathbb{N}$.

3.2 Event-relativized Formulas

In the following, we identify a class of formulas for which the point-based and the interval-based semantics coincide. For formulas in this class, a point-based monitoring algorithm can be used to soundly monitor properties given by formulas interpreted using the interval-based semantics. We assume that the propositions are typed, i.e., $P = S \cup E$, where S contains the state propositions and E the event propositions, and a family of signals $\hat{\gamma} = (\gamma_p)_{p \in S \cup E}$ is consistent with S and E . Moreover, we assume without loss of generality that there is always at least one event signal γ in $\hat{\gamma}$ that is the infinite union of singular intervals, e.g., γ is the signal of a clock event that regularly occurs over time.

We inductively define the sets rel_{\forall} and rel_{\exists} for formulas in negation normal form. Recall that a formula is in negation normal form if negation only occurs directly in front of propositions. A logically-equivalent negation normal form of a formula can always be obtained by eliminating double negations and by pushing negations inwards, where we consider the Boolean connective \vee and the temporal operator “trigger” \top_I as primitives. Note that $\phi \top_I \psi = \neg(\neg\phi \text{S}_I \neg\psi)$.

$$\neg p \in rel_{\forall} \quad \text{if} \quad p \in E \quad (\forall 1)$$

$$\phi_1 \vee \phi_2 \in rel_{\forall} \quad \text{if} \quad \phi_1 \in rel_{\forall} \text{ or } \phi_2 \in rel_{\forall} \quad (\forall 2)$$

$$\phi_1 \wedge \phi_2 \in rel_{\forall} \quad \text{if} \quad \phi_1 \in rel_{\forall} \text{ and } \phi_2 \in rel_{\forall} \quad (\forall 3)$$

$$p \in rel_{\exists} \quad \text{if} \quad p \in E \quad (\exists 1)$$

$$\phi_1 \wedge \phi_2 \in rel_{\exists} \quad \text{if} \quad \phi_1 \in rel_{\exists} \text{ or } \phi_2 \in rel_{\exists} \quad (\exists 2)$$

$$\phi_1 \vee \phi_2 \in rel_{\exists} \quad \text{if} \quad \phi_1 \in rel_{\exists} \text{ and } \phi_2 \in rel_{\exists} \quad (\exists 3)$$

A formula ϕ is *event-relativized* if $\alpha \in rel_{\forall}$ and $\beta \in rel_{\exists}$, for every subformula of ϕ of the form $\alpha \text{S}_I \beta$ or $\beta \top_I \alpha$. We call the formula ϕ *strongly* event-relativized if ϕ is event-relativized and $\phi \in rel_{\forall} \cup rel_{\exists}$.

The following theorem relates the interval-based semantics and the point-based semantics for event-relativized formulas.

Theorem 1. *Let $\hat{\gamma} = (\gamma_p)_{p \in S \cup E}$ be a family of consistent signals and $\bar{\tau}$ the time sequence listing the occurrences of events in $\hat{\gamma}$, i.e., $\bar{\tau}$ is the time sequence obtained by linearly ordering the set $\bigcup_{p \in E} \gamma_p$. For an event-relativized formula ϕ and every $i \in \mathbb{N}$, it holds that*

$$\hat{\gamma}, \tau_i \models \phi \quad \text{iff} \quad \hat{\gamma}, \bar{\tau}, i \models \phi.$$

Furthermore, if ϕ is strongly event-relativized, then it also holds that (a) $\hat{\gamma}, \tau \not\models \phi$ if $\phi \in rel_{\exists}$ and (b) $\hat{\gamma}, \tau \models \phi$ if $\phi \in rel_{\forall}$, for all $\tau \in \mathbb{T} \setminus \{\tau_i \mid i \in \mathbb{N}\}$.

Observe that the formulas in Example 1 and 2 are not event-relativized. The definition of event-relativized formulas and Theorem 1 straightforwardly extend to richer real-time logics that also contain future operators and are first-order. We point out that most formulas that we encountered when formalizing security policies in such a richer temporal logic are strongly event-relativized [3].

From Theorem 1, it follows that the interval-based semantics can simulate the point-based one by using a fresh event proposition sp with its signal $\gamma_{sp} = \{\tau_i \mid i \in \mathbb{N}\}$, for a time sequence $\bar{\tau}$. We then event-relativize a formula ϕ with the proposition sp , i.e., subformulas of the form $\psi_1 \mathbf{S}_I \psi_2$ are replaced by $(sp \rightarrow \psi_1) \mathbf{S}_I (sp \wedge \psi_2)$ and $\psi_1 \mathbf{T}_I \psi_2$ by $(sp \wedge \psi_1) \mathbf{T}_I (sp \rightarrow \psi_2)$.

4 Monitoring Algorithms

In this section, we present and analyze our monitoring algorithms for both the point-based and the interval-based setting. Without loss of generality, the algorithms assume that the temporal subformulas of a formula ϕ occur only once in ϕ . Moreover, let P be the set of propositions that occur in ϕ .

4.1 A Point-based Monitoring Algorithm

Our monitoring algorithm for the point-based semantics iteratively computes the truth values of a formula ϕ at the sample-points $i \in \mathbb{N}$ for a given time sequence $\bar{\tau}$ and a family of signals $\hat{\gamma} = (\gamma_p)_{p \in P}$. We point out that $\bar{\tau}$ and $\hat{\gamma}$ are given incrementally, i.e., in the $(i+1)$ st iteration, the monitor obtains the time-stamp τ_i and the signals between the previous time-stamp and τ_i . In fact, in the point-based setting, we do not need to consider “chunks” of signals; instead, we can restrict ourselves to the snapshots $\Gamma_i := \{p \in P \mid \tau_i \in \gamma_p\}$, for $i \in \mathbb{N}$, i.e., Γ_i is the set of propositions that hold at time τ_i .

Each iteration of the monitor is performed by executing the procedure **step**[•]. At sample-point $i \in \mathbb{N}$, **step**[•] takes as arguments the formula ϕ , the snapshot Γ_i , and i 's time-stamp τ_i . It computes the truth value of ϕ at i recursively over ϕ 's structure. For efficiency, the procedure **step**[•] maintains for each subformula ψ of the form $\psi_1 \mathbf{S}_I \psi_2$ a sequence L_ψ of time-stamps. These sequences are initialized by the procedure **init**[•] and updated by the procedure **update**[•]. These three procedures² are given in Figure 2 and are described next.

The base case of **step**[•] where ϕ is a proposition and the cases for the Boolean connectives \neg and \wedge are straightforward. The only involved case is where ϕ is of the form $\phi_1 \mathbf{S}_I \phi_2$. In this case, **step**[•] first updates the sequence L_ϕ and then computes ϕ 's truth value at the sample-point $i \in \mathbb{N}$.

Before we describe how we update the sequence L_ϕ , we describe the elements that are stored in L_ϕ and how we obtain from them ϕ 's truth value. After the update of L_ϕ by **update**[•], the sequence L_ϕ stores the time-stamps τ_j with $\tau_i - \tau_j \in \leq I$ (i.e., the time-stamps that satisfy the time constraint now or that might satisfy it in the future) at which ϕ_2 holds and from which ϕ_1 continuously holds up to the current sample-point i (i.e., ϕ_2 holds at $j \leq i$ and ϕ_1 holds at each $k \in \{j+1, \dots, i\}$). Moreover, if there are time-stamps τ_j and $\tau_{j'}$ with $j < j'$ in L_ϕ with $\tau_i - \tau_j \in I$ and $\tau_i - \tau_{j'} \in I$ then we only keep in L_ϕ the time-stamp

² Our pseudo-code is written in a functional-programming style using pattern matching. $\langle \rangle$ denotes the empty sequence, $++$ sequence concatenation, and $x :: L$ the sequence with head x and tail L .

```

step•( $\phi, \Gamma, \tau$ )
  case  $\phi = p$ 
    return  $p \in \Gamma$ 
  case  $\phi = \neg\phi'$ 
    return not step•( $\phi', \Gamma, \tau$ )
  case  $\phi = \phi_1 \wedge \phi_2$ 
    return step•( $\phi_1, \Gamma, \tau$ ) and step•( $\phi_2, \Gamma, \tau$ )
  case  $\phi = \phi_1 \text{ S}_I \phi_2$ 
    update•( $\phi, \Gamma, \tau$ )
    if  $L_\phi = \langle \rangle$  then return false
    else return  $\tau - \text{head}(L_\phi) \in I$ 

init•( $\phi$ )
  for each  $\psi \in \text{sf}(\phi)$  with  $\psi = \psi_1 \text{ S}_I \psi_2$  do
     $L_\psi := \langle \rangle$ 

update•( $\phi, \Gamma, \tau$ )
  let  $\phi_1 \text{ S}_I \phi_2 = \phi$ 
   $b_1 = \text{step}^{\bullet}(\phi_1, \Gamma, \tau)$ 
   $b_2 = \text{step}^{\bullet}(\phi_2, \Gamma, \tau)$ 
   $L = \text{if } b_1 \text{ then drop}^{\bullet}(L_\phi, I, \tau) \text{ else } \langle \rangle$ 
  in if  $b_2$  then  $L_\phi := L \uparrow \langle \tau \rangle$ 
  else  $L_\phi := L$ 

```

Fig. 2. Monitoring in a point-based setting.

```

drop•( $L, I, \tau$ )
  case  $L = \langle \rangle$ 
    return  $\langle \rangle$ 
  case  $L = \kappa :: L'$ 
    if  $\tau - \kappa \not\leq I$  then return drop•( $L', I, \tau$ )
    else return drop•( $\kappa, L', I, \tau$ )

drop•( $\kappa, L', I, \tau$ )
  case  $L' = \langle \rangle$ 
    return  $\langle \kappa \rangle$ 
  case  $L' = \kappa' :: L''$ 
    if  $\tau - \kappa' \in I$  then return drop•( $\kappa', L'', I, \tau$ )
    else return  $\kappa :: L'$ 

```

Fig. 3. Auxiliary procedures.

of the later sample-point, i.e., τ_j . Finally, the time-stamps in L_ϕ are ordered increasingly. Having L_ϕ at hand, it is easy to determine ϕ 's truth value. If L_ϕ is the empty sequence then obviously ϕ does not hold at sample-point i . If L_ϕ is non-empty then ϕ holds at i iff the first time-stamp κ in L_ϕ fulfills the timing constraints given by the interval I , i.e., $\tau_i - \kappa \in I$. Recall that ϕ holds at i iff there is a sample-point $j \leq i$ with $\tau_i - \tau_j \in I$ at which ϕ_2 holds and since then ϕ_1 continuously holds.

Initially, L_ϕ is the empty sequence. If ϕ_2 holds at sample-point i , then **update**[•] adds the time-stamp τ_i to L_ϕ . However, prior to this, it removes the time-stamps of the sample-points from which ϕ_1 does not continuously hold. Clearly, if ϕ_1 does not hold at i then we can empty the sequence L_ϕ . Otherwise, if ϕ_1 holds at i , we first drop the time-stamps for which the distance to the current time-stamp τ_i became too large with respect to the right margin of I . Afterwards, we drop time-stamps until we find the last time-stamp τ_j with $\tau_i - \tau_j \in I$. This is done by the procedures **drop**[•] and **drop**'[•] shown in Figure 3.

Theorem 2. *Let ϕ be a formula, $\hat{\gamma} = (\gamma_p)_{p \in P}$ be a family of signals, $\bar{\tau}$ be a time sequence, and $n > 0$. The procedure **step**[•]($\phi, \Gamma_{n-1}, \tau_{n-1}$) terminates, and returns **true** iff $\hat{\gamma}, \bar{\tau}, n - 1 \models \phi$, whenever **init**[•](ϕ), **step**[•](ϕ, Γ_0, τ_0), \dots , **step**[•]($\phi, \Gamma_{n-2}, \tau_{n-2}$) were called previously in this order, where $\Gamma_i = \{p \in P \mid \tau_i \in \gamma_p\}$, for $i < n$.*

We end this subsection by analyzing the monitor's computational complexity. Observe that we cannot bound the space that is needed to represent the time-stamps in the time sequence $\bar{\tau}$. They become arbitrarily large as time progresses. Moreover, since the time domain is dense, they can be arbitrarily close to each

other. As a consequence, operations like subtraction of elements from \mathbb{T} cannot be done in constant time. We return to this point in Section 4.3.

In the following, we assume that each $\tau \in \mathbb{T}$ is represented by two bit strings for the numerator and denominator. The representation of an interval I consists of the representations for $\ell(I)$ and $r(I)$ and whether the left margin and right margin is closed or open. We denote the maximum length of these bit strings by $\|\tau\|$ and $\|I\|$, respectively. The operations on elements in \mathbb{T} that the monitoring algorithm performs are subtractions and membership tests. Subtraction $\tau - \tau'$ can be carried out in time $\mathcal{O}(m^2)$, where $m = \max\{\|\tau\|, \|\tau'\|\}$.³ A membership test $\tau \in I$ can also be carried out in time $\mathcal{O}(m^2)$, where $m = \max\{\|\tau\|, \|I\|\}$.

The following theorem establishes an upper bound on the time complexity of our monitoring algorithm.

Theorem 3. *Let $\phi, \hat{\gamma}, \bar{\tau}, n$, and $\Gamma_0, \dots, \Gamma_{n-1}$ be as in Theorem 2. Executing the sequence $\text{init}^\bullet(\phi), \text{step}^\bullet(\phi, \Gamma_0, \tau_0), \dots, \text{step}^\bullet(\phi, \Gamma_{n-1}, \tau_{n-1})$ requires $\mathcal{O}(m^2 \cdot n \cdot |\phi|)$ time, where $m = \max(\{\|I\| \mid \alpha \mathbf{S}_I \beta \in \text{sf}(\phi)\} \cup \{\|\tau_0\|, \dots, \|\tau_{n-1}\|\})$.*

4.2 An Interval-based Monitoring Algorithm

Our monitoring algorithm for the interval-based semantics determines, for a given family of signals $\hat{\gamma} = (\gamma_p)_{p \in P}$, the truth value of a formula ϕ , for any $\tau \in \mathbb{T}$. In other words, it determines the set $\gamma_{\phi, \hat{\gamma}} := \{\tau \in \mathbb{T} \mid \hat{\gamma}, \tau \models \phi\}$. We simply write γ_ϕ instead of $\gamma_{\phi, \hat{\gamma}}$ when the family of signals $\hat{\gamma}$ is clear from the context. Similar to the point-based setting, the monitor incrementally receives the input $\hat{\gamma}$ and incrementally outputs γ_ϕ , i.e., the input and output signals are split into “chunks” by an infinite interval partition \bar{J} . Concretely, the input of the $(i+1)$ st iteration consists of the formula ϕ that is monitored, the interval J_i of \bar{J} , and the family $\hat{\Delta}_i = (\Delta_{i,p})_{p \in P}$ of sequences of intervals $\Delta_{i,p} = \overline{\text{ip}}^1(\gamma_p \cap J_i)$, for propositions $p \in P$. The output of the $(i+1)$ st iteration is the sequence $\overline{\text{ip}}^1(\gamma_\phi \cap J_i)$.

Observe that the sequence $\overline{\text{ip}}^1(\gamma_p \cap J_i)$ only consists of a finite number of intervals since the signal γ_p satisfies the finite-variability condition and J_i is bounded. Moreover, since γ_p is stable on every interval in $\overline{\text{ip}}(\gamma_p)$ and an interval has a finite representation, the sequence $\overline{\text{ip}}^1(\gamma_p \cap J_i)$ finitely represents the signal chunk $\gamma_p \cap J_i$. Similar observations are valid for the signal chunk $\gamma_\phi \cap J_i$.

Each iteration is performed by the procedure **step**. To handle the since operator efficiently, **step** maintains for each subformula ψ of the form $\psi_1 \mathbf{S}_I \psi_2$, a (possibly empty) interval K_ψ and a finite sequence of intervals Δ_ψ . These global variables are initialized by the procedure **init** and updated by the procedure **update**. These three procedures are given in Figure 4 and are described next.

The procedure **step** computes the signal chunk $\gamma_\phi \cap J_i$ recursively over the formula structure. It utilizes the right-hand sides of the following equalities:

$$\gamma_p \cap J_i = \bigcup_{K \in \overline{\text{ip}}^1(\gamma_p \cap J_i)} K \quad (1)$$

³ Note that $\frac{p}{q} - \frac{p'}{q'} = \frac{p \cdot q' - p' \cdot q}{q \cdot q'}$ and that $\mathcal{O}(m^2)$ is an upper bound on the multiplication of two m bit integers. There are more sophisticated algorithms for multiplication that run in $\mathcal{O}(m \log m \log \log m)$ time [19] and $\mathcal{O}(m \log m 2^{\log^* m})$ time [8]. For simplicity, we use the quadratic upper bound.

```

step( $\phi, \hat{\Delta}, J$ )
  case  $\phi = p$ 
    return  $\Delta_p$ 
  case  $\phi = \neg\phi'$ 
    let  $\Delta' = \text{step}(\phi', \hat{\Delta}, J)$ 
    in return invert( $\Delta', J$ )
  case  $\phi = \phi_1 \wedge \phi_2$ 
    let  $\Delta_1 = \text{step}(\phi_1, \hat{\Delta}, J)$ 
         $\Delta_2 = \text{step}(\phi_2, \hat{\Delta}, J)$ 
    in return intersect( $\Delta_1, \Delta_2$ )
  case  $\phi = \phi_1 S_I \phi_2$ 
    let ( $\Delta'_1, \Delta'_2$ ) = update( $\phi, \hat{\Delta}, J$ )
    in return merge(combine( $\Delta'_1, \Delta'_2, I, J$ ))

init( $\phi$ )
  for each  $\psi \in \text{sf}(\phi)$  with  $\psi = \psi_1 S_I \psi_2$  do
     $K_\psi := \emptyset$ 
     $\Delta_\psi := \langle \rangle$ 

update( $\phi, \hat{\Delta}, J$ )
  let  $\phi_1 S_I \phi_2 = \phi$ 
     $\Delta_1 = \text{step}(\phi_1, \hat{\Delta}, J)$ 
     $\Delta_2 = \text{step}(\phi_2, \hat{\Delta}, J)$ 
     $\Delta'_1 = \text{prepend}(K_\phi, \Delta_1)$ 
     $\Delta'_2 = \text{concat}(\Delta_\phi, \Delta_2)$ 
  in  $K_\phi := \text{if } \Delta'_1 = \langle \rangle \text{ then } \emptyset \text{ else last}(\Delta'_1)$ 
     $\Delta_\phi := \text{drop}(\Delta'_2, I, J)$ 
  return ( $\Delta'_1, \Delta'_2$ )

```

Fig. 4. Monitoring in an interval-based setting

```

cons( $K, \Delta$ )
  if  $K = \emptyset$  then
    return  $\Delta$ 
  else
    return  $K :: \Delta$ 

invert( $\Delta, J$ )
  case  $\Delta = \langle \rangle$ 
    return  $\langle J \rangle$ 
  case  $\Delta = K :: \Delta'$ 
    return cons( $J \cap <^<K, \text{invert}(\Delta', J \cap (K^>))$ )

intersect( $\Delta_1, \Delta_2$ )
  if  $\Delta_1 = \langle \rangle$  or  $\Delta_2 = \langle \rangle$  then
    return  $\langle \rangle$ 
  else
    let  $K_1 :: \Delta'_1 = \Delta_1$ 
         $K_2 :: \Delta'_2 = \Delta_2$ 
    in if  $K_1 \cap (K_2^>) = \emptyset$  then
      return cons( $K_1 \cap K_2, \text{intersect}(\Delta'_1, \Delta'_2)$ )
    else
      return cons( $K_1 \cap K_2, \text{intersect}(\Delta_1, \Delta'_2)$ )

```

Fig. 5. The auxiliary procedures for the Boolean connectives.

$$\gamma_{\neg\phi'} \cap J_i = J_i \setminus \left(\bigcup_{K \in \overline{\text{ip}}^1(\gamma_{\phi'} \cap J_i)} K \right) \quad (2)$$

$$\gamma_{\phi_1 \wedge \phi_2} \cap J_i = \bigcup_{\substack{K_1 \in \overline{\text{ip}}^1(\gamma_{\phi_1} \cap J_i) \\ K_2 \in \overline{\text{ip}}^1(\gamma_{\phi_2} \cap J_i)}} (K_1 \cap K_2) \quad (3)$$

$$\gamma_{\phi_1 S_I \phi_2} \cap J_i = \bigcup_{\substack{K_1 \in \overline{\text{ip}}^1(\gamma_{\phi_1}) \text{ with } K_1 \cap J_i \neq \emptyset \\ K_2 \in \overline{\text{ip}}^1(\gamma_{\phi_2}) \text{ with } (K_2 \oplus I) \cap (J_i^>) \neq \emptyset}} \left(((K_2 \cap {}^+K_1) \oplus I) \cap K_1 \cap J_i \right) \quad (4)$$

where ${}^+K := \{\ell(K)\} \cup K$, for $K \in \mathbb{I}$, i.e., making the interval K left-closed.

The equalities (1), (2), and (3) are obvious and their right-hand sides are directly reflected in our pseudo-code. The case where ϕ is a proposition is straightforward. For the case $\phi = \neg\phi'$, we use the procedure `invert`, shown in Figure 5, to compute $\overline{\text{ip}}^1(\gamma_\phi \cap J_i)$ from $\Delta' = \overline{\text{ip}}^1(\gamma_{\phi'} \cap J_i)$. This is done by “complementing” Δ' with respect to the interval J_i . For instance, the output of `invert`($\langle [1, 2] (3, 4), [0, 10] \rangle$) is $\langle [0, 1] (2, 3) [4, 10] \rangle$. For the case $\phi = \phi_1 \wedge \phi_2$, we use the procedure `intersect`, also shown in Figure 5, to compute $\overline{\text{ip}}^1(\gamma_\phi \cap J_i)$ from $\Delta_1 = \overline{\text{ip}}^1(\gamma_{\phi_1} \cap J_i)$ and $\Delta_2 = \overline{\text{ip}}^1(\gamma_{\phi_2} \cap J_i)$. This procedure returns the sequence of intervals that have a non-empty intersection of two intervals in the input sequences. The elements in the returned sequence are ordered increasingly.

The equality (4) for $\phi = \phi_1 S_I \phi_2$ is less obvious and using its right-hand side for an implementation is also less straightforward since the intervals K_1 and K_2

```

prepend( $K, \Delta$ )
  if  $K = \emptyset$  then
    return  $\Delta$ 
  else
    case  $\Delta = \langle \rangle$ 
      return  $\langle K \rangle$ 
    case  $\Delta = K' :: \Delta'$ 
      if adjacent( $K, K'$ ) or  $K \cap K' \neq \emptyset$  then
        return  $K \cup K' :: \Delta'$ 
      else
        return  $K :: \Delta$ 

concat( $\Delta_1, \Delta_2$ )
  case  $\Delta_1 = \langle \rangle$ 
    return  $\Delta_2$ 
  case  $\Delta_1 = \Delta'_1 ++ \langle K_1 \rangle$ 
    return  $\Delta'_1 ++ \text{prepend}(K_1, \Delta_2)$ 

drop( $\Delta'_2, I, J$ )
  case  $\Delta'_2 = \langle \rangle$ 
    return  $\langle \rangle$ 
  case  $\Delta'_2 = K_2 :: \Delta''_2$ 
    let  $K = (K_2 \oplus I) \cap (J^>)$ 
    in if  $K = \emptyset$  then return drop( $\Delta''_2, I, J$ )
       else return drop'( $K, \Delta'_2, I, J$ )

combine( $\Delta'_1, \Delta'_2, I, J$ )
  if  $\Delta'_1 = \langle \rangle$  or  $\Delta'_2 = \langle \rangle$  then return  $\langle \rangle$ 
  else
    let  $K_2 :: \Delta''_2 = \Delta'_2$ 
    in if  $(K_2 \oplus I) \cap J = \emptyset$  then return  $\langle \rangle$ 
       else
         let  $K_1 :: \Delta'_1 = \Delta'_1$ 
          $\Delta = \text{if } K_2^> \cap {}^+K_1 = \emptyset \text{ then}$ 
           combine( $\Delta'_1, \Delta'_2, I, J$ )
         else
           combine( $\Delta'_1, \Delta''_2, I, J$ )
         in return  $(K_2 \cap {}^+K_1) \oplus I) \cap K_1 \cap J :: \Delta$ 

merge( $\Delta$ )
  case  $\Delta = \langle \rangle$ 
    return  $\Delta$ 
  case  $\Delta = K :: \Delta'$ 
    return prepend( $K, \text{merge}(\Delta')$ )

drop'( $K, \Delta'_2, I, J$ )
  case  $\Delta'_2 = \langle \rangle$ 
    return  $\langle K \rangle$ 
  case  $\Delta'_2 = K_2 :: \Delta''_2$ 
    let  $K' = (K_2 \oplus I) \cap (J^>)$ 
    in if  $K \subseteq K'$  then return drop'( $K', \Delta''_2, I, J$ )
       else return  $\Delta'_2$ 

```

Fig. 6. The auxiliary procedures for the since operator.

are not restricted to occur in the current chunk J_i . Instead, they are intervals in $\overline{\text{ip}}^1(\gamma_{\phi_1})$ and $\overline{\text{ip}}^1(\gamma_{\phi_2})$, respectively, with certain constraints.

Before giving further implementation details, we first show why equality (4) holds. To prove the inclusion \subseteq , assume $\tau \in \gamma_{\phi_1} \mathcal{S}_I \phi_2 \cap J_i$. By the semantics of the since operator, there is a $\tau_2 \in \gamma_{\phi_2}$ with $\tau - \tau_2 \in I$ and $\tau_1 \in \gamma_{\phi_1}$, for all $\tau_1 \in (\tau_2, \tau]$.

- Obviously, $\tau_2 \in K_2$, for some $K_2 \in \overline{\text{ip}}^1(\gamma_{\phi_2})$. By taking the time constraint I into account, K_2 satisfies the constraint $(K_2 \oplus I) \cap (J_i^>) \neq \emptyset$. Note that even the more restrictive constraint $(K_2 \oplus I) \cap J_i \neq \emptyset$ holds. However, we employ the weaker constraint in our implementation as it is useful for later iterations.
- Since $\overline{\text{ip}}(\gamma_{\phi_1})$ is the coarsest interval partition of γ_{ϕ_1} , there is an interval $K_1 \in \overline{\text{ip}}^1(\gamma_{\phi_1})$ with $(\tau_2, \tau] \subseteq K_1$. As $\tau \in J_i$, the constraint $K_1 \cap J_i \neq \emptyset$ holds.

It follows that $\tau \in K_1$ and $\tau_2 \in {}^+K_1$, and thus $\tau_2 \in K_2 \cap {}^+K_1$. From $\tau - \tau_2 \in I$, we obtain that $\tau \in (K_2 \cap {}^+K_1) \oplus I$. Finally, since $\tau \in K_1 \cap J_i$, we have that $\tau \in ((K_2 \cap {}^+K_1) \oplus I) \cap K_1 \cap J_i$. The other inclusion \supseteq can be shown similarly.

For computing the signal chunk $\gamma_{\phi_1} \mathcal{S}_I \phi_2 \cap J_i$, the procedure **step** first determines the subsequences Δ'_1 and Δ'_2 of $\overline{\text{ip}}^1(\gamma_{\phi_1})$ and $\overline{\text{ip}}^1(\gamma_{\phi_2})$ consisting of those intervals K_1 and K_2 appearing in the equality (4), respectively. This is done by the procedure **update**. Afterwards, **step** computes the sequence $\overline{\text{ip}}^1(\gamma_{\phi} \cap J_i)$ from Δ'_1 and Δ'_2 by using the procedures **combine** and **merge**, given in Figure 6. We now explain how $\text{merge}(\text{combine}(\Delta'_1, \Delta'_2, I, J))$ returns the sequence $\overline{\text{ip}}^1(\gamma_{\phi_1} \mathcal{S}_I \phi_2 \cap J_i)$. First, $\text{combine}(\Delta'_1, \Delta'_2, I, J)$ computes a sequence of intervals

whose union is $\gamma_{\phi_1} s_I \phi_2 \cap J_i$. It traverses the ordered sequences Δ'_1 and Δ'_2 and adds the interval $((K_2 \cap {}^+K_1) \oplus I) \cap K_1 \cap J_i$ to the resulting ordered sequence, for K_1 in Δ'_1 and K_2 in Δ'_2 . The test $K_2 \cap {}^+K_1 = \emptyset$ determines in which sequence (Δ'_1 or Δ'_2) we advance next: if the test succeeds then $K'_2 \cap {}^+K_1 = \emptyset$ where K'_2 is the successor of K_2 in Δ'_2 , and hence we advance in Δ'_1 . The sequence Δ'_2 is not necessarily entirely traversed: when $(K_2 \oplus I) \cap J_i = \emptyset$, one need not inspect other elements K'_2 of the sequence Δ'_2 , as then $((K'_2 \cap {}^+K_1) \oplus I) \cap K_1 \cap J_i = \emptyset$. The elements in the sequence returned by the `combine` procedure might be empty, adjacent, or overlapping. The `merge` procedure removes empty elements and merges adjacent or overlapping intervals, i.e., it returns the sequence $\overline{\text{pp}}^1(\gamma_{\phi_1} s_I \phi_2 \cap J_i)$.

Finally, we explain the contents of the variables K_ϕ and Δ_ϕ and how they are updated. We start with K_ϕ . At the $(i+1)$ st iteration, for some $i \geq 0$, the following invariant is satisfied by K_ϕ : before the update, the interval K_ϕ is the last interval of $\overline{\text{pp}}^1(\gamma_{\phi_1} \cap \leq J_{i-1})$ if $i > 0$ and this sequence is not empty, and K_ϕ is the empty set otherwise. The interval K_ϕ is prepended to the sequence $\overline{\text{pp}}^1(\gamma_{\phi_1} \cap J_i)$ using the `prepend` procedure from Figure 6, which merges K_ϕ with the first interval of $\Delta_1 = \overline{\text{pp}}^1(\gamma_{\phi_1} \cap J_i)$ if these two intervals are adjacent. The obtained sequence Δ'_1 is the maximal subsequence of $\overline{\text{pp}}^1(\gamma_{\phi_1} \cap \leq J_i)$ such that $K_1 \cap J_i \neq \emptyset$, for each interval K_1 in Δ'_1 . Thus, after the update, K_ϕ is the last interval of $\overline{\text{pp}}^1(\gamma_{\phi_1} \cap \leq J_i)$ if this sequence is not empty, and K_ϕ is the empty set otherwise. Hence the invariant on K_ϕ is preserved at the next iteration.

The following invariant is satisfied by Δ_ϕ at the $(i+1)$ st iteration: before the update, the sequence Δ_ϕ is empty if $i = 0$, and otherwise, if $i > 0$, it stores the intervals K_2 in $\overline{\text{pp}}^1(\gamma_{\phi_2} \cap \leq J_{i-1})$ with $(K_2 \oplus I) \cap (J_{i-1}^>) \neq \emptyset$ and $(K_2 \oplus I) \cap (J_{i-1}^>) \not\subseteq (K'_2 \oplus I) \cap (J_{i-1}^>)$, where K'_2 is the successor of K_2 in $\overline{\text{pp}}^1(\gamma_{\phi_2} \cap \leq J_{i-1})$. The procedure `concat` concatenates the sequence Δ_ϕ with the sequence $\Delta_2 = \overline{\text{pp}}^1(\gamma_{\phi_2} \cap J_i)$. Since the last interval of Δ_ϕ and the first interval of Δ_2 can be adjacent, `concat` might need to merge them. Thus, the obtained sequence Δ'_2 is a subsequence of $\overline{\text{pp}}^1(\gamma_{\phi_2} \cap \leq J_i)$ such that $(K_2 \oplus I) \cap (J_i^>) \neq \emptyset$, for each element K_2 . Note that $J_{i-1}^> = J_i^>$. The updated sequence Δ_ϕ is obtained from Δ'_2 by removing the intervals K_2 with $(K_2 \oplus I) \cap (J_i^>) = \emptyset$, i.e., the intervals that are irrelevant for later iterations. The procedure `drop` from Figure 6 removes these intervals. Moreover, if there are intervals K_2 and K'_2 in Δ_ϕ with $(K_2 \oplus I) \cap (J_i^>) \subseteq (K'_2 \oplus I) \cap (J_i^>)$ then only the interval that occurs later is kept in Δ_ϕ . This is done by the procedure `drop'`. Thus, after the update, the sequence Δ_ϕ stores the intervals K_2 in $\overline{\text{pp}}^1(\gamma_{\phi_2} \cap \leq J_i)$ with $(K_2 \oplus I) \cap (J_i^>) \neq \emptyset$ and $(K_2 \oplus I) \cap (J_i^>) \not\subseteq (K'_2 \oplus I) \cap (J_i^>)$, where K'_2 is the successor of K_2 in $\overline{\text{pp}}^1(\gamma_{\phi_2} \cap \leq J_i)$. Hence the invariant on Δ_ϕ is preserved at the next iteration.

Theorem 4. *Let ϕ be a formula, $\hat{\gamma} = (\gamma_p)_{p \in P}$ a family of signals, \bar{J} an infinite interval partition, and $n > 0$. The procedure `step`($\phi, \hat{\Delta}_{n-1}, J_{n-1}$) terminates and returns the sequence $\overline{\text{pp}}^1(\gamma_\phi \cap J_{n-1})$, whenever `init`(ϕ), `step`($\phi, \hat{\Delta}_0, J_0$), \dots , `step`($\phi, \hat{\Delta}_{n-2}, J_{n-2}$) were called previously in this order, where $\hat{\Delta}_i = (\Delta_{i,p})_{p \in P}$ with $\Delta_{i,p} = \overline{\text{pp}}^1(\gamma_p \cap J_i)$, for $i < n$.*

Finally, we analyze the monitor’s computational complexity. As in the point-based setting, we take the representation size of elements of the time domain \mathbb{T} into account. The basic operations here in which elements of \mathbb{T} are involved are operations on intervals like checking emptiness (i.e. $I = \emptyset$), “extension” (e.g. $I^>$), and “shifting” (i.e. $I \oplus J$). The representation size of the interval $I \oplus J$ is in $\mathcal{O}(\|I\| + \|J\|)$. The time to carry out the shift operation is in $\mathcal{O}(\max\{\|I\|, \|J\|\}^2)$. All the other basic operations that return an interval do not increase the representation size of the resulting interval with respect to the given intervals. However, the time complexity is quadratic in the representation size of the given intervals whenever the operation needs to compare interval margins.

The following theorem establishes an upper bound on the time complexity of our monitoring algorithm.

Theorem 5. *Let ϕ , $\hat{\gamma}$, \bar{J} , n , and $\hat{\Delta}_i$ be given as in Theorem 4. Executing the sequence $\text{init}(\phi)$, $\text{step}(\phi, \hat{\Delta}_0, J_0)$, \dots , $\text{step}(\phi, \hat{\Delta}_{n-1}, J_{n-1})$ requires $\mathcal{O}(m^2 \cdot (n + \delta \cdot |\phi|) \cdot |\phi|^3)$ time, where $m = \max(\{\|I\| \mid \alpha \mathbf{S}_I \beta \in \text{sf}(\phi)\} \cup \{\|J_0\|, \dots, \|J_{n-1}\|\}) \cup \bigcup_{p \in P} \{\|K\| \mid K \in \overline{\text{pp}}^1(\gamma_p \cap (<J_n))\})$ and $\delta = \sum_{p \in P} \|\gamma_p \cap (<J_n)\|$.*

We remark that the factor $m^2 \cdot |\phi|^2$ is due to the operations on the margins of intervals. With the assumption that the representation of elements of the time domain is constant, we obtain the upper bound $\mathcal{O}((n + \delta \cdot |\phi|) \cdot |\phi|)$.

4.3 Time Domains

The stated worst-case complexities of both monitoring algorithms take the representation size of the elements in the time domain into account. In practice, it is often reasonable to assume that these elements have a bounded representation, since arbitrarily precise clocks do not exist. For example, for many applications it suffices to represent time-stamps as Unix time, i.e., 32 or 64 bit signed integers. The operations performed by our monitoring algorithms on the time domain elements would then be carried out in constant time. However, a consequence of this practically motivated assumption is that the time domain is discrete and bounded rather than dense and unbounded.

For a discrete time domain, we must slightly modify the interval-based monitoring algorithm, namely, the operator ^+K used in the equality (4) must be redefined. In a discrete time domain, we extend K by one point in time to the left if it exists, i.e., $^+K := K \cup \{k - 1 \mid k \in K \text{ and } k > 0\}$. No modifications are needed for the point-based algorithm. If we assume a discrete and unbounded time domain, we still cannot assume that the operations on elements from the time domain can be carried out in constant time. But multiplication is no longer needed to compare elements in the time domain and thus the operations can be carried in time linear in the representation size. The worst-case complexity of both algorithms improves accordingly.

When assuming limited-precision clocks, which results in a discrete time domain, a so-called fictitious-clock semantics [2, 18] is often used. This semantics formalizes, for example, that if the system event e happens strictly before the

event e' but both events fall between two clock ticks, then we can distinguish them by temporal ordering, not by time. In a fictitious-clock semantics, we time-stamp e and e' with the same clock value and in a trace e appears strictly before e' . For ordering e and e' in a trace, signals must be synchronized. Our point-based monitoring algorithm can directly be used for a fictitious-clock semantics. It iteratively processes a sequence of snapshots $\langle \Gamma_0, \Gamma_1, \dots \rangle$ together with a sequence of time-stamps $\langle \tau_0, \tau_1, \dots \rangle$, which is increasing but not necessarily strictly increasing anymore. In contrast, our interval-based monitoring algorithm does not directly carry over to a fictitious-clock semantics.

4.4 Comparison of the Monitoring Algorithms

In the following, we compare our two algorithms when monitoring a strongly event-relativized formula ϕ . By Theorem 1, the point-based setting and the interval-based setting coincide on this formula class.

First note that the input for the $(i+1)$ th iteration of the point-based monitoring algorithm can be easily obtained online from the given signals $\hat{\gamma} = (\gamma)_{p \in S \cup E}$. Whenever an event occurs, we record the time $\tau_i \in \mathbb{T}$, determine the current truth values of the propositions, i.e., $\Gamma_i = \{p \in P \mid \tau_i \in \gamma_p\}$, and invoke the monitor by executing $\text{step}^\bullet(\phi, \Gamma_i, \tau_i)$. The worst-case complexity of the point-based monitoring algorithm of the first n iterations is $\mathcal{O}(m^2 \cdot n \cdot |\phi|)$, where m is according to Theorem 3.

When using the interval-based monitoring algorithm, we are more flexible in that we need not invoke the monitoring algorithm whenever an event occurs. Instead, we can freely split the signals into chunks. Let \bar{J} be a splitting in which the n' th interval $J_{n'-1}$ is right-closed and $r(J_{n'-1}) = \tau_{n-1}$. We have the worst-case complexity of $\mathcal{O}(m'^2 \cdot (n' + \delta \cdot |\phi|) \cdot |\phi|^3)$, where m' and δ are according to Theorem 5. We can lower this upper bound, since the formula ϕ is strongly event-relativized. Instead of the factor $m'^2 \cdot |\phi|^2$ for processing the interval margins in the n' iterations, we only have the factor m'^2 . The reason is that the margins of the intervals in the signal chunks of subformulas of the form $\psi_1 S_I \psi_2$ already appear as interval margins in the input.

Note that $m' \geq m$ and that δ is independent of n' . Under the assumption that $m' = m$, the upper bounds on the running times for different splittings only differ by n' , i.e., how often we invoke the procedure step . The case where $n' = 1$ corresponds to the scenario where we use the monitoring algorithm offline (up to time τ_{n-1}). The case where $n' = n$ corresponds to the case where we invoke the monitor whenever an event occurs. Even when using the interval-based monitoring algorithm offline and assuming constant representation of the elements in \mathbb{T} , the upper bounds differ by the factors n and $\delta \cdot |\phi|$. Since $\delta \geq n$, the upper bound of the point-based monitoring algorithm is lower. In fact, the examples in Appendix C show that the gap between the running times matches our upper bounds and that $\delta \cdot |\phi|$ can be significantly larger than n .

5 Related Work

We only discuss the monitoring algorithms most closely related to ours, namely, those of Basin et al. [4], Thati and Roşu [20], and Nickovic and Maler [14, 15].

The point-based monitoring algorithms here simplify and optimize the monitoring algorithm of Basin et al. [4] given for the future-bounded fragment of metric first-order temporal logic. We restricted ourselves here to the propositional setting and to the past-only fragment of metric temporal logic to compare the effect of different time models on monitoring.

Thati and Roşu [20] provide a monitoring algorithm for metric temporal logic with a point-based semantics, which uses formula rewriting. Their algorithm is more general than ours for the point-based setting since it handles past and future operators. Their complexity analysis is based on the assumption that operations involving elements from the time domain can be carried out in constant time. The worst-case complexity of their algorithm on the past-only fragment is worse than ours, since rewriting a formula can generate additional formulas. In particular, their algorithm is not linear in the number of subformulas.

Nickovic and Maler’s [14, 15] monitoring algorithms are for the interval-based setting and have ingredients similar to our algorithm for this setting. These ingredients were first presented by Nickovic and Maler for an offline version of their monitoring algorithms [13] for the fragment of interval metric temporal logic with bounded future operators. Their setting is more general in that their signals are continuous functions and not Boolean values for each point in time. Moreover, their algorithms also handle bounded [15] and unbounded [14] future operators by delaying the evaluation of subformulas. The algorithm in [14] slightly differs from the one in [15]: [14] also handles past operators and before starting monitoring, it rewrites the given formula to eliminate the temporal operators until and since with timing constraints. The main difference to our algorithm is that Maler and Nickovic do not provide algorithmic details for handling the Boolean connectives and the temporal operators. In fact, the worst-case complexity, which is only stated for their offline algorithm [13], seems to be too low even when ignoring representation and complexity issues for elements of the time domain.

We are not aware of any work that compares different time models for runtime verification. The surveys [2, 6, 16] on real-time logics focus on expressiveness, satisfiability, and automatic verification of real-time systems. A comparison of a point-based and interval-based time model for temporal databases with a discrete time domain is given by Toman [21]. The work by Furia and Rossi [9] on sampling and the work on digitization [11] by Henzinger et al. are orthogonal to our comparison. These relate fragments of metric interval temporal logic with respect to a discrete and a dense time domain.

6 Conclusions

We have presented, analyzed, and compared monitoring algorithms for real-time logics with point-based and interval-based semantics. Our comparison provides a

detailed explanation of trade-offs between the different time models with respect to monitoring. Moreover, we have presented a practically relevant fragment for the interval-based setting by distinguishing between state variables and system events, which can be more efficiently monitored in the point-based setting.

As future work, we plan to extend the monitoring algorithms to handle bounded future operators. This includes analyzing their computational complexities and comparing them experimentally. Another line of research is to establish lower bounds for monitoring real-time logics. Thati and Roşu [20] give lower bounds for future fragments of metric temporal logic including the next operator. However, we are not aware of any lower bounds for the past-only fragment.

References

1. R. Alur, T. Feder, and T. A. Henzinger. The benefits of relaxing punctuality. *J. ACM*, 43(1):116–146, 1996.
2. R. Alur and T. A. Henzinger. Logics and models of real time: A survey. In *Proceedings of the 1991 REX Workshop on Real-Time: Theory in Practice*, volume 600 of *Lect. Notes Comput. Sci.*, pages 74–106. Springer, 1992.
3. D. Basin, F. Klaedtke, and S. Müller. Monitoring security policies with metric first-order temporal logic. In *Proceedings of the 15th ACM Symposium on Access Control Models and Technologies (SACMAT)*, pages 23–33. ACM Press, 2010.
4. D. Basin, F. Klaedtke, S. Müller, and B. Pfitzmann. Runtime monitoring of metric first-order temporal properties. In *Proceedings of the 28th IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'08)*, volume 2 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 49–60. Schloss Dagstuhl - Leibniz Center for Informatics, 2008.
5. A. Bauer, M. Leucker, and C. Schallhart. Monitoring of real-time properties. In *Proceedings of the 26th International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 4337 of *Lect. Notes Comput. Sci.*, pages 260–272. Springer, 2006.
6. P. Bouyer. Model-checking times temporal logics. In *Proceedings of the 5th Workshop on Methods for Modalities (M4M5)*, volume 231 of *Elec. Notes Theo. Comput. Sci.*, pages 323–341. Elsevier Science Inc., 2009.
7. D. Drusinsky. On-line monitoring of metric temporal logic with time-series constraints using alternating finite automata. *J. UCS*, 12(5):482–498, 2006.
8. M. Fürer. Faster integer multiplication. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC)*, pages 55–67. ACM Press, 2007.
9. C. A. Furia and M. Rossi. A theory of sampling for continuous-time metric temporal logic. *ACM Trans. Comput. Log.*, 12(1), 2010.
10. A. Goodloe and L. Pike. Monitoring distributed real-time systems: A survey and future directions. Technical Report NASA/CR-2010-216724, NASA Langley Research Center, July 2010.
11. T. A. Henzinger, Z. Manna, and A. Pnueli. What good are digital clocks? In *In Proceedings of the 19th International Colloquium on Automata, Languages and Programming (ICALP)*, volume 623 of *Lect. Notes Comput. Sci.*, pages 545–558. Springer, 1992.
12. K. J. Kristoffersen, C. Pedersen, and H. R. Andersen. Runtime verification of timed LTL using disjunctive normalized equation systems. In *Proceedings of the 3rd*

- Workshop on Runtime Verification (RV)*, volume 89 of *Elec. Notes Theo. Comput. Sci.*, pages 210–225. Elsevier Science Inc., 2003.
13. O. Maler and D. Nickovic. Monitoring temporal properties of continuous signals. In *Proceedings of the Joint International Conferences on Formal Modelling and Analysis of Timed Systems (FORMATS) and on Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT)*, volume 3253 of *Lect. Notes Comput. Sci.*, pages 152–166. Springer, 2004.
 14. D. Ničković. *Checking Timed and Hybrid Properties: Theory and Applications*. PhD thesis, Université Joseph Fourier, Grenoble, France, October 2008.
 15. D. Nickovic and O. Maler. AMT: A property-based monitoring tool for analog systems. In *Proceedings of the 5th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS)*, volume 4763 of *Lect. Notes Comput. Sci.*, pages 304–319. Springer, 2007.
 16. J. Ouaknine and J. Worrell. Some recent results in metric temporal logic. In *Proceedings of the 6th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS)*, volume 5215 of *Lect. Notes Comput. Sci.*, pages 1–13. Springer, 2008.
 17. L. Pike, A. Goodloe, R. Morisset, and S. Niller. Copilot: A hard real-time runtime monitor. In *Proceedings of the 1st International Conference on Runtime Verification (RV)*, volume 6418 of *Lect. Notes Comput. Sci.*, pages 345–359. Springer, 2010.
 18. J.-F. Raskin and P.-Y. Schobbens. Real-time logics: Fictitious clock as an abstraction of dense time. In *Proceedings of the 3rd International Workshop on Tools and Algorithms for Construction and Analysis of Systems (TACAS)*, volume 1217 of *Lect. Notes Comput. Sci.*, pages 165–182. Springer, 1997.
 19. A. Schönhage and V. Strassen. Schnelle Multiplikation großer Zahlen. *Computing*, 7(3–4):281–292, 1971.
 20. P. Thati and G. Roşu. Monitoring algorithms for metric temporal logic specifications. In *Proceedings of the 4th Workshop on Runtime Verification (RV)*, volume 113 of *Elec. Notes Theo. Comput. Sci.*, pages 145–162. Elsevier Science Inc., 2005.
 21. D. Toman. Point vs. interval-based query languages for temporal databases (extended abstract). In *Proceedings of the 15th ACM Symposium on Principles of Database Systems (PODS)*, pages 58–67. ACM Press, 1996.

A Proof Details on the Event-relativized Fragment

We prove Theorem 1 by induction on the structure of the formula ϕ . Let T be the set $\mathbb{T} \setminus \{\tau_i \mid i \in \mathbb{N}\}$.

Base case: $\phi = p$ with $p \in P$. If p is a state proposition then there is nothing to prove. Assume that p is an event proposition. By definition, p is strongly event-relativized, in particular, $p \in rel_{\exists}$. In the interval-based semantics for $\tau \in \mathbb{T}$, it holds that $\hat{\gamma}, \tau \models p$ iff $\tau \in \gamma_p$. Since p is an event proposition, $\tau = \tau_i$, for some $i \in \mathbb{N}$. It follows that $\hat{\gamma}, \tau_i \models p$ iff $\gamma_p, \bar{\tau}, i \dot{\models} p$. Note that $\hat{\gamma}, \tau \not\models p$ when $\tau \in T$.

Base case: $\phi = \neg p$ with $p \in P$. If p is a state proposition then there is nothing to prove. Assume that p is an event proposition. By definition, $\neg p$ is strongly event-relativized, in particular, $\neg p \in rel_{\forall}$. In the interval-based semantics for $\tau \in \mathbb{T}$, it holds that $\hat{\gamma}, \tau \models \neg p$ iff $\tau \notin \gamma_p$. Since p is an event proposition, if $\tau = \tau_i$ then $\tau_i \notin \gamma_p$, for all $i \in \mathbb{N}$. It follows that $\hat{\gamma}, \tau_i \models \neg p$ iff $\gamma_p, \bar{\tau}, i \dot{\models} \neg p$. Note that $\hat{\gamma}, \tau \not\models \neg p$ when $\tau \in T$.

Step case: $\phi = \phi_1 \wedge \phi_2$. We have the following equivalences: $\hat{\gamma}, \tau_i \models \phi_1 \wedge \phi_2$ iff (by the interval-based semantics) $\hat{\gamma}, \tau_i \models \phi_1$ and $\hat{\gamma}, \tau_i \models \phi_2$ iff (by the induction hypothesis) $\hat{\gamma}, \bar{\tau}, i \dot{\models} \phi_1$ and $\hat{\gamma}, \bar{\tau}, i \dot{\models} \phi_2$ iff (by the point-based semantics) $\hat{\gamma}, \bar{\tau}, i \dot{\models} \phi_1 \wedge \phi_2$.

If $\phi \in rel_{\exists}$ then by definition, $\phi_1 \in rel_{\exists}$ or $\phi_2 \in rel_{\exists}$. Without loss of generality, assume $\phi_1 \in rel_{\exists}$. By the induction hypothesis, $\hat{\gamma}, \tau \not\models \phi_1$, for all $\tau \in T$. It follows that $\hat{\gamma}, \tau \not\models \phi$, for all $\tau \in T$.

If $\phi \in rel_{\forall}$ then by definition, $\phi_1 \in rel_{\exists}$ and $\phi_2 \in rel_{\exists}$. By the induction hypothesis, $\hat{\gamma}, \tau \not\models \phi_1$, for all $\tau \in T$ and $\hat{\gamma}, \tau \not\models \phi_2$, for all $\tau \in T$. It follows that $\hat{\gamma}, \tau \not\models \phi$, for all $\tau \in T$.

Step case: $\phi = \phi_1 \vee \phi_2$. This case is dual to the previous case. We omit it.

Step case: $\phi = \phi_1 \mathbf{S}_I \phi_2$. Since ϕ is not strongly event-relativized, we need only show that $\hat{\gamma}, \tau_i \models \phi$ iff $\hat{\gamma}, \bar{\tau}, i \dot{\models} \phi$. We have the equivalence

$$\hat{\gamma}, \tau_i \models \phi_1 \mathbf{S}_I \phi_2 \quad \text{iff} \quad \begin{array}{l} \text{there is some } \tau \in [0, \tau_i) \text{ with } \tau_i - \tau \in I, \\ \hat{\gamma}, \tau \models \phi_2, \text{ and } \hat{\gamma}, \kappa \models \phi_2, \text{ for all } \kappa \in (\tau, \tau_i]. \end{array}$$

Since $\phi_2 \in rel_{\exists}$, there is some $j \in \mathbb{N}$ with $\tau_j = \tau$. From the induction hypothesis and the fact that $\phi_1 \in rel_{\forall}$, we conclude that

$$\hat{\gamma}, \tau_i \models \phi_1 \mathbf{S}_I \phi_2 \quad \text{iff} \quad \begin{array}{l} \text{there is some } j \leq i \text{ with } \tau_i - \tau_j \in I, \\ \hat{\gamma}, \bar{\tau}, j \dot{\models} \phi_2, \text{ and } \hat{\gamma}, \bar{\tau}, k \dot{\models} \phi_1, \text{ for all } k \in \{j+1, \dots, i\}. \end{array}$$

In the point-based semantics, the right-hand side is by definition equivalent to $\hat{\gamma}, \bar{\tau}, i \dot{\models} \phi_1 \mathbf{S}_I \phi_2$.

Step case: $\phi = \phi_1 \mathbf{T}_I \phi_2$. This case is dual to the previous case. We omit it.

B Proof Details on Complexity Analysis

In the following, $\mathbf{tsf}(\phi)$ denotes the set of the temporal subformulas of ϕ , i.e., $\mathbf{tsf}(\phi) := \{\psi \in \mathbf{sf}(\phi) \mid \psi \text{ is of the form } \psi_1 \mathbf{S}_I \psi_2\}$, and $\mathbf{dsf}(\phi)$ denotes the direct

subformulas of ϕ , i.e. $\text{dsf}(p) = \emptyset$, for $p \in P$, $\text{dsf}(\neg\phi) = \{\phi\}$, and $\text{dsf}(\phi_1 \wedge \phi_2) = \text{dsf}(\phi_1) \cup \text{dsf}(\phi_2) = \{\phi_1, \phi_2\}$. We say that a formula is a temporal formula if it is of the form $\alpha S_I \beta$. Note that $p \wedge \blacklozenge q$ is not considered a temporal formula.

B.1 Point-based Monitoring Algorithm

For proving Theorem 3, we first analyze the running time of a single iteration. We claim that the running time of $\text{step}^\bullet(\phi, \Gamma_{n-1}, \tau_{n-1})$ is in $\mathcal{O}(|\phi| + m^2 \cdot \sum_{\psi \in \text{tsf}(\phi)} T_\psi^n)$, where $T_\psi^1 := 1$ and

$$T_{\alpha S_I \beta}^n := 1 + |\{\tau_j \mid \tau_{n-2} - \ell(I) \leq \tau_j \leq \tau_{n-1} - \ell(I), \text{ for some } j < n\}|,$$

for $n > 1$.

Since we traverse ϕ 's syntax tree recursively, the running time of one iteration is the sum of all t_ψ , for all occurrences of subformulas ψ of ϕ , where t_ψ denotes the running time for ψ without the running times for its proper subformulas.

We have that $t_\psi \in \mathcal{O}(1)$ for the cases where ψ is of the form p , $\neg\psi'$, or $\psi_1 \wedge \psi_2$. Note that we can assume, without loss of generality, that the membership test $p \in \Gamma$ for the base case in the procedure step^\bullet can be done in constant time. The reason is that, in the n th iteration, from the set Γ_n (which is given for instance as a list) we can first build a hash table that allows us to check in constant time whether the proposition p is an element of Γ_n . Building (and discarding) such a hash table takes $\mathcal{O}(|P|)$ time. Since $|P| \leq |\phi|$, this term is subsumed by the claimed complexity $\mathcal{O}(|\phi| + m^2 \sum_{\psi \in \text{tsf}(\phi)} T_\psi^n)$.

It remains to analyze the running time for the case where ψ is of the form $\psi_1 S_I \psi_2$. We first make the following observations about the sequence L_ψ and the elements it contains in the n th iteration.

- (1) For each element τ in L_ψ , we have that $\|\tau\| \leq m$, since τ is a time-stamp that occurs in the prefix of length n of the time sequence $\bar{\tau}$.
- (2) Removing the head and appending an element to L_ψ can be done in $\mathcal{O}(m)$, since we assume that L_ψ is implemented as a doubly linked list with pointers to the first and to the last element.
- (3) The disequality test $\text{last}(L) \neq \tau$ and the membership tests whether the distance $\tau - \kappa$ is in I or $\leq I$ can be performed in time $\mathcal{O}(m^2)$, since the time-stamps τ and κ occur in the prefix of the time sequence $\bar{\tau}$, and thus, by assumption, $\|I\|, \|\tau\|, \|\kappa\| \leq m$.

From these observations, it is easy to see $t_\psi \in \mathcal{O}(m^2 \cdot T)$, where T is number of elements from the sequence L_ψ that are visited by the procedures drop^\bullet and drop'^\bullet . Note that L_ψ is empty in the first iteration. Suppose that $n > 1$. The procedure drop^\bullet first traverses the sequence L_ψ up to the first element τ_k such that $\tau_{n-1} - \tau_k \notin I$. Hence all elements τ_j in L_ψ up to and excluding τ_k satisfy $\tau_j \leq \tau_{n-1} - \ell(I)$. Moreover, except for at most the first element of L_ψ , all elements τ_j in L_ψ satisfy $\tau_j \geq \tau_{n-2} - \ell(I)$. It follows that $T \leq 1 + T_\psi^n$, since at most two elements (the first one and the last one visited in L_ψ) may be outside the interval $[\tau_{n-2} - \ell(I), \tau_{n-1} - \ell(I)]$.

We conclude that $\text{step}^\bullet(\phi, I_{n-1}, \tau_{n-1})$ has the claimed the running time $\mathcal{O}(|\phi| + m^2 \cdot \sum_{\psi \in \text{tsf}(\phi)} T_\psi^n)$.

It remains to prove the upper bound on the running time of all n iterations, i.e., for the sequence $\text{init}^\bullet(\phi)$, $\text{step}^\bullet(\phi, I_0, \tau_0)$, \dots , $\text{step}^\bullet(\phi, I_{n-1}, \tau_{n-1})$. To establish this upper bound, note that the sets $\{\tau_j \mid \tau_{i-1} - \ell(I) \leq \tau_j \leq \tau_i - \ell(I), \text{ for some } j \leq i\}$ and $\{\tau_j \mid \tau_i - \ell(I) \leq \tau_j \leq \tau_{i+1} - \ell(I), \text{ for some } j \leq i+1\}$ have at most one element in common. Thus, $\sum_{1 < i \leq n} T_\psi^i \leq n + (n + |\{\tau_j \mid j < n\}|)$. Since $|\{\tau_j \mid j < n\}| \leq n$, this sum of the T_ψ^i s is in $\mathcal{O}(n)$. Then, by summing up the running times of all iterations, we obtain that the total running time is in $\mathcal{O}(n \cdot |\phi| + m^2 \cdot n \cdot |\text{tsf}(\psi)|)$, from which the stated upper bound follows.

B.2 Interval-based Monitoring Algorithm

In the following, we prove Theorem 5. Let ψ be a subformula of ϕ . We denote by $t_n(\psi)$ the running time of the sequence $\text{init}(\psi)$, $\text{step}(\psi, \hat{\Delta}_0, J_0)$, \dots , $\text{step}(\psi, \hat{\Delta}_{n-1}, J_{n-1})$.

We also define m_ψ as

$$\max \{ \|I\| \mid \alpha \text{S}_I \beta \in \text{sf}(\psi) \} \cup \{ \|J_0\|, \dots, \|J_{n-1}\| \} \cup \bigcup_{\psi' \in \text{sf}(\psi)} \{ \|K\| \mid K \in \text{ip}^1(\gamma_{\psi'} \cap \prec J_n) \},$$

i.e., m_ψ is the maximal representation size of some interval that occurs in the first n iterations of the monitoring algorithm when determining the signal for ψ . By inspecting the operations that are performed on intervals in one iteration, we obtain

$$m_\psi \leq \begin{cases} m_{\psi'} & \text{if } \psi = \neg\psi', \\ \max\{m_{\psi_1}, m_{\psi_2}\} & \text{if } \psi = \psi_1 \wedge \psi_2, \\ \max\{m_{\psi_1}, m_{\psi_2}\} + \|I\| & \text{if } \psi = \psi_1 \text{S}_I \psi_2, \\ m & \text{otherwise.} \end{cases}$$

From these inequalities, by induction on the formula structure, the inequality $m_\psi \leq m \cdot |\psi|$ straightforwardly follows.

The following lemma, which follows from the equalities (2)–(4), establishes an upper bound on the number of intervals that are necessary for representing the signal determined by the formula ψ up to some point in time.

Lemma 1. *Let J be a bounded interval with $0 \in J$. If ψ is not a proposition then*

$$\|\gamma_\psi \cap J\| \in \mathcal{O}\left(\sum_{\psi' \in \text{dsf}(\psi)} \|\gamma_{\psi'} \cap J\|\right).$$

The following lemma is the key ingredient for proving Theorem 5. It establishes an upper bound of the running time for the formula ψ excluding the running times for its proper subformulas.

Lemma 2. *If ψ is a proposition then $t_n(\psi) \in \mathcal{O}(n)$. If ψ is not a proposition then*

$$t_n(\psi) - \sum_{\psi' \in \text{dsf}(\psi)} t_n(\psi') \in \mathcal{O}\left(m_\psi^2 \cdot \sum_{\psi' \in \text{dsf}(\psi)} \sum_{i < n} (1 + \|\gamma_{\psi'} \cap J_i\|)\right).$$

Proof. The upper bound when ψ is a proposition is obviously true, since in each iteration we just return Δ_p . In the following we prove the upper bound for the cases when ψ is not a proposition. We do this by case distinction.

In the following, $|\Delta|$ denotes the length of a sequence Δ .

Case $\psi = \neg\psi'$. For $i < n$, the running time of $\text{step}(\psi, \hat{\Delta}_i, J_i)$ without the running time of $\text{step}(\psi', \hat{\Delta}_i, J_i)$ equals the running time of $\text{invert}(\Delta', J_i)$. The procedure invert visits sequentially the elements in the sequence Δ' . Each visit costs $\mathcal{O}(m_\psi^2)$, since the procedure cons checks whether an interval K is empty before appending it to the inverted sequence. From Theorem 4 we know that at the $(i + 1)$ st iteration, the sequence Δ' equals $\overline{\text{pp}}^1(\gamma_{\psi'} \cap J_i)$. So, the length of Δ' is $|\Delta'| = \|\gamma_{\psi'} \cap J_i\|$. From this follows the upper bound for $\neg\psi'$.

Case $\psi = \psi_1 \wedge \psi_2$. Similar arguments as in the case where ψ equals $\neg\psi'$ establish the upper bound for $\psi = \psi_1 \wedge \psi_2$. Note that $\text{intersect}(\Delta_1, \Delta_2)$ runs in time $\mathcal{O}(m_\psi^2 \cdot (1 + \|\gamma_{\psi_1} \cap J_i\| + \|\gamma_{\psi_2} \cap J_i\|))$.

Case $\psi = \psi_1 \text{ S}_I \psi_2$. We first inspect the running time of the $(i + 1)$ st iteration with $i < n$. The running time of $\text{step}(\psi, \hat{\Delta}_i, J_i)$ without the running times for $\text{step}(\psi_1, \hat{\Delta}_i, J_i)$ and $\text{step}(\psi_2, \hat{\Delta}_i, J_i)$ equals the sum of the running times of $\text{prepend}(K_\phi, \Delta_1)$, $\text{concat}(\Delta_\phi, \Delta_2)$, $\text{last}(\Delta'_1)$, $\text{drop}(\Delta'_2, I, J_i)$, which are called by the procedure update , and of $\text{merge}(\text{combine}(\Delta'_1, \Delta'_2, I, J_i))$.

The procedures prepend , concat , and last run in time at most $\mathcal{O}(m_\psi^2)$. So, their total running time for all n iterations is in $\mathcal{O}(m_\psi^2 \cdot n)$. The procedures merge , combine , and drop , including the call to drop' , visit sequentially elements of input sequence. Each such visit costs $\mathcal{O}(m_\psi^2)$. Whereas the procedure merge visits all elements, the procedures combine , drop and drop' might stop before reaching the end of the input sequence.

Before analyzing the procedures drop , combine , and merge , we make the following remarks. Consider the procedure update . We have $|\Delta'_1| \leq 1 + |\Delta_1|$ and $|\Delta'_2| \leq |\Delta_\phi| + |\Delta_2|$, where $|\Delta_\phi|$ refers to the number of elements in Δ_ϕ before calling the procedure drop . Moreover, from Theorem 4, $|\Delta'_1| \in \mathcal{O}(\|\gamma_{\psi_1} \cap J_i\|)$ and $\Delta'_2 \subseteq \overline{\text{pp}}^1(\gamma_{\psi_2} \cap \leq J_i)$.

We first focus on the call to $\text{drop}(\Delta'_2, I, J_i)$. The drop procedure only visits a prefix of the sequence Δ'_2 , and returns the last visited element appended to the unvisited suffix of Δ'_2 . Thus the running time of drop is in $\mathcal{O}(m_\psi^2 \cdot (1 + |\Delta_{diff}|))$, where Δ'_ϕ is the value of Δ_ϕ after the call to drop and Δ_{diff} is such that $\Delta'_2 = \Delta_{diff} ++ \Delta'_\phi$. In the next iteration, the input of drop will be a subsequence of $\overline{\text{pp}}^1(\gamma_{\psi_2} \cap \leq J_{i+1})$ having Δ'_ϕ as a prefix. It follows that at most one element in $\overline{\text{pp}}^1(\gamma_{\psi_2} \cap \leq J_n)$ is visited by any two consecutive calls to the drop procedure. Thus the total running time of the drop procedure during the first n iterations is in $\mathcal{O}(m_\psi^2 \cdot (n + \|\gamma_{\psi_2} \cap \leq J_n\|))$. Note that $\leq J_n = \leq J_{n-1}$.

We now focus on the running time of $\text{combine}(\Delta'_1, \Delta'_2, I, J_i)$. The combine procedure traverses the sequences Δ'_1 and Δ'_2 , and, similarly to the procedure intersect , it runs in $\mathcal{O}(m_\psi^2 \cdot (1 + |\Delta'_1| + |\Delta'_2|))$. This upper bound does not suffice to obtain the desired upper bound on the total running time, as we do not have that $|\Delta'_2| \in \mathcal{O}(\|\gamma_{\psi_2} \cap J_i\|)$. Nevertheless, we prove below that, while the sequences Δ'_2

in two consecutive iterations may have more than one interval in common, at most one such interval is visited by both iterations. This shows that in the first n iterations at most $n + \|\gamma_{\psi_2} \cap \prec J_n\|$ intervals from the sequences Δ'_2 are visited.

To show that at most one interval is visited by the `combine` procedure during each of the $(i+1)$ st and $(i+2)$ nd iterations, we recall that in the $(i+1)$ st iteration, we have $\Delta'_2 = \Delta_{diff} \uparrow \Delta'_\phi$, while in the $(i+2)$ nd iteration the sequence Δ'_ϕ is a prefix of Δ'_2 . Hence it is sufficient to show that at most one interval among the ones that `combine` visits in the $(i+1)$ st iteration belongs to Δ'_ϕ . Let $used(\Delta'_2) := \{K_2 \in \Delta'_2 \mid (K_2 \oplus I) \cap J_i \neq \emptyset\}$. Note that there is at most one visited element in Δ'_2 that is not in $used(\Delta'_2)$, as in the $(i+1)$ st iteration, the `combine` procedure stops whenever $(K_2 \oplus I) \cap J_i = \emptyset$ and thus no elements following K_2 in Δ'_2 are visited. We show that there is at most one interval in both Δ'_ϕ and $used(\Delta'_2)$. Suppose by absurdity that there are two such intervals. Then there are two consecutive such intervals, K_2 and K'_2 . By the invariant on Δ'_ϕ enforced by the procedure `drop'`, we have $(K_2 \oplus I) \cap (J_i^\succ) \not\subseteq (K'_2 \oplus I) \cap (J_i^\succ)$. Then $K'_2 \oplus I \subseteq J_i^\succ$, and thus $(K'_2 \oplus I) \cap J_i = \emptyset$, which is a contradiction with $K'_2 \in used(\Delta'_2)$.

Finally, the running time of the procedure `merge` is linear in length of its argument, that is, the length of the result of `combine`, which in turn is at most $|\Delta'_1| + |\Delta'_2|$. From the above discussion, it follows that the total running time of `merge` for the n iterations is in $\mathcal{O}(m_\psi^2 \cdot (n + \|\gamma_{\psi_2} \cap \prec J_n\| + \sum_{i < n} (1 + \|\gamma_{\psi_1} \cap J_i\|)))$. As $\|\gamma_{\psi_2} \cap \prec J_n\| \leq \sum_{i < n} \|\gamma_{\psi_2} \cap J_i\|$, this concludes the proof of the upper bound, where ψ is of the form $\psi_1 \mathbf{S}_I \psi_2$. \square

We now put the ingredients together to establish the upper bound of Theorem 5. By Lemma 2, we have that

$$t_n(\psi) - \sum_{\psi' \in \text{dsf}(\psi)} t_n(\psi') \in \mathcal{O}\left(m_\psi^2 \cdot \sum_{\psi' \in \text{dsf}(\psi)} \sum_{i < n} (1 + \|\gamma_{\psi'} \cap J_i\|)\right).$$

We obtain

$$t_n(\psi) - \sum_{\psi' \in \text{dsf}(\psi)} t_n(\psi') \in \mathcal{O}\left(m_\psi^2 \cdot \sum_{\psi' \in \text{dsf}(\psi)} (n + \|\gamma_{\psi'} \cap \prec J_n\|)\right),$$

since $\sum_{i < n} (1 + \|\gamma \cap J_i\|) \in \mathcal{O}(n + \|\gamma \cap \prec J_n\|)$, for any signal $\gamma \subseteq \mathbb{T}$. Since ψ has at most two direct subformulas, we obtain

$$t_n(\psi) - \sum_{\psi' \in \text{dsf}(\psi)} t_n(\psi') \in \mathcal{O}\left(m_\psi^2 \cdot (n + \sum_{\psi' \in \text{dsf}(\psi)} \|\gamma_{\psi'} \cap \prec J_n\|)\right).$$

With Lemma 1, we get

$$t_n(\psi) - \sum_{\psi' \in \text{dsf}(\psi)} t_n(\psi') \in \mathcal{O}\left(m_\psi^2 \cdot (n + |\psi| \cdot \sum_{p \in P} \|\gamma_p \cap \prec J_n\|)\right),$$

that is, $t_n(\psi) - \sum_{\psi' \in \text{dsf}(\psi)} t_n(\psi') \in \mathcal{O}(m_\psi^2 \cdot (n + |\psi| \cdot \delta))$. Using the inequalities $m_\psi \leq m \cdot |\psi|$ and $|\psi| \leq |\phi|$, we obtain

$$t_n(\psi) - \sum_{\psi' \in \text{dsf}(\psi)} t_n(\psi') \in \mathcal{O}(m^2 \cdot |\phi|^2 \cdot (n + |\phi| \cdot \delta)).$$

Summing up these equations for each subformula ψ of ϕ , we get

$$t_n(\phi) \in \mathcal{O}(n \cdot |\phi| + m^2 \cdot |\phi|^2 \cdot (n + |\phi| \cdot \delta) \cdot |\phi|),$$

which simplifies to $t_n(\phi) \in \mathcal{O}(m^2 \cdot (n + |\phi| \cdot \delta) \cdot |\phi|^3)$.

C Worst-case Examples

We present instances illustrating the gap between the computational complexity of the monitoring algorithms. Recall that when using the interval-based monitoring algorithm offline and assuming constant representation of the elements in \mathbb{T} , the upper bounds differ by the factors n and $\delta \cdot |\phi|$, for the point-based and the interval-based algorithm, respectively. In the following examples, we assume that time-stamps and interval margins are represented within constant space.

Our first example shows that δ can be significantly larger than n . It also illustrates that the point-based algorithm ignores large portions of the state signals, while the interval-based algorithm processes the whole state signals.

Example 3. Let ϕ be the formula $e \wedge s$, where e is an event proposition and s is a state proposition. Note that ϕ is a strongly event-relativized. Let γ_e be the event signal \mathbb{N} and γ_s be the state signal $\bigcup_{i \in \mathbb{N}} [2i \cdot \frac{n}{2k}, (2i+1) \cdot \frac{n}{2k})$, for some $n, k \in \mathbb{N}$ with $n, k > 0$. The running time of the first n iterations of the point-based algorithm is clearly in $\Theta(n)$. The running time of the first iteration of the interval-based algorithm on the chunk $\gamma_s \cap [0, n]$ is in $\Theta(k)$ since the procedure `intersect` traverses the entire sequence $\overline{\text{ip}}^1(\gamma_s \cap [0, n])$, which contains k intervals. The choices of n and k are independent to each other. When choosing k significantly larger than n , the number of intervals in $\overline{\text{ip}}^1(\gamma_s \cap [0, n])$ dominates the running time of the interval-based algorithm, and not the number of events, which in turn determine the running time of the point-based algorithm.

Our second example shows that $\Omega(|\phi|^2)$ is a lower bound for the worst-case running time of the interval-based algorithm, even when the proposition set P is a singleton. It again illustrates that for event-relativized formulas, the interval-based algorithm processes for each subformula intermediate signal chunks that contain information also on what happens when events do not occur, while this information is (soundly) ignored by the point-based algorithm.

Example 4. Let p be an event proposition. We define the event-relativized formulas $\phi_1 := \blacklozenge_{[1,1]} p$, and $\phi_i := \phi_{i-1} \vee \blacklozenge_{[i,i]} p$, for $i \in \mathbb{N}$ with $i > 0$. In the following, let $k \in \mathbb{N}$ with $k > 0$ and let ϕ be the formula $p \wedge \phi_k$. Note that ϕ has the form $p \wedge ((\blacklozenge_{[1,1]} p) \vee (\blacklozenge_{[2,2]} p) \vee \dots \vee (\blacklozenge_{[k,k]} p))$ and that $|\phi| \in \Theta(k)$. Consider the event signal $\gamma_p = [0, 0]$. We have that $\gamma_\phi = \emptyset$ and $\gamma_{\phi_i} = [1, 1] \cup [2, 2] \cup \dots \cup [i, i]$, for $i \in \mathbb{N}$ with $1 \leq i \leq k$. Let \bar{J} be an interval partition with $r(J_0) > k$. We have $\|\overline{\text{ip}}^1(\gamma_p \cap J_0)\| = 1$ and $\|\overline{\text{ip}}^1(\gamma_{\phi_i} \cap J_0)\| = i$ and $\|\overline{\text{ip}}^1(\gamma_{\neg\phi_i} \cap J_0)\| = i + 1$, for each $i \in \mathbb{N}$ with $1 \leq i \leq k$. The running time of the first iteration of the interval-based algorithm is in $\Theta(k^2)$, thus in $\Theta(|\phi|^2)$, as for each subformula ϕ_i , which is $\neg(\neg\phi_{i-1} \wedge \neg\blacklozenge_{[i,i]} p)$ when removing the syntactic sugar for the Boolean connective \vee , the procedures `invert` and `intersect` traverse $\Theta(i)$ intervals, for $1 < i \leq k$. The running time of the point-based algorithm on the signal chunk given by J_0 is in $\Theta(|\phi|)$, since there is only one event, namely, the one that occurs at time 0.

Note that by considering disjunction as a primitive and adjusting the implementation accordingly would not change the bound $\Theta(|\phi|^2)$. Finally, we remark

that singular intervals attached to temporal operators do not play an essential role in this example: we obtain the same running times when replacing the intervals $[i, i]$ with intervals $[i - \epsilon, i + \epsilon]$, where $\epsilon \in \mathbb{T}$ is sufficiently small.