

Robust Online Monitoring of Signal Temporal Logic

Jyotirmoy V. Deshmukh¹, Alexandre Donzé², Shromona Ghosh²
Xiaoqing Jin¹, Garvit Juniwal², and Sanjit A. Seshia²

¹ Toyota Technical Center, `firstname.lastname@tema.toyota.com`

² University of California Berkeley,
{donze, shromona.ghosh, garvitjuniwal, sseschia}@eecs.berkeley.edu

Abstract. Signal Temporal Logic (STL) is a formalism used to rigorously specify requirements of cyberphysical systems (CPS), i.e., systems mixing digital or discrete components in interaction with a continuous environment or analog components. STL is naturally equipped with a quantitative semantics which can be used for various purposes: from assessing the robustness of a specification to guiding searches over the input and parameter space with the goal of falsifying the given property over system behaviors. Algorithms have been proposed and implemented for *offline* computation of such quantitative semantics, but only few methods exist for an *online* setting, where one would want to monitor the satisfaction of a formula during simulation. In this paper, we formalize a semantics for robust online monitoring of *partial* traces, i.e., traces for which there might not be enough data to decide the Boolean satisfaction (and to compute its quantitative counterpart). We propose an efficient algorithm to compute it and demonstrate its usage on two large scale real-world case studies coming from the automotive domain and from CPS education in a Massively Open Online Course (MOOC) setting. We show that savings in computationally expensive simulations far outweigh any overheads incurred by an online approach.

1 Introduction

Design engineers for embedded control software typically validate their designs by inspecting concrete observations of system behavior. For instance, in the model-based development (MBD) paradigm, designers have access to numerical simulation tools to obtain traces from models of systems. An important problem is then to be able to efficiently test whether some logical property φ holds for a given simulation trace. It is increasingly common [13, 9, 12, 2, 15] to specify such properties using a real-time temporal logic such as Signal Temporal Logic (STL) [7] or Metric Temporal Logic (MTL) [10]. An *offline monitoring* approach involves performing an *a posteriori* analysis on *complete* simulation traces (i.e., traces starting at time 0, and lasting till a user-specified time horizon). Theoretical and practical results for offline monitoring [10, 5, 7, 17] focus on the efficiency of monitoring as a function of the length of the trace, and the size of the formula representing the property φ .

There are a number of situations where offline monitoring is unsuitable. Consider the case where the monitor is to be deployed in an actual system to detect erroneous behavior. As embedded software is typically resource constrained, offline monitoring – which requires storing the entire observed trace – is impractical. Also, when a monitor

is used in a simulation-based validation tool, a single simulation may run for several minutes or even hours. If we wish to monitor a safety property over the simulation, a better use of resources is to abort the simulation whenever a violation is detected. Such situations demand an *online monitoring algorithm*, which has markedly different requirements. In particular, a good online monitoring algorithm must: (1) be able to generate intermediate estimates of property satisfaction based on *partial signals*, (2) use minimal amount of data storage, and (3) be able to run fast enough in a real-time setting.

Most works on online monitoring algorithms for logics such as Linear Temporal Logic (LTL) or Metric Temporal Logic (MTL) have focussed on the Boolean satisfaction of properties by partial signals [11, 8, 18]. However, recent work has shown that by assigning quantitative semantics to real-time logics such as MTL and STL, problems such as bug-finding, parameter synthesis, and robustness analysis can be solved using powerful off-the-shelf optimization tools [1, 4]. A robust satisfaction value is a function mapping a property φ and a trace $\mathbf{x}(t)$ to a real number. A large positive value suggests that $\mathbf{x}(t)$ easily satisfies φ , a positive value close to zero suggests that $\mathbf{x}(t)$ is close to violating φ , and a negative value indicates a violation of φ . While the recursive definitions of quantitative semantics naturally define offline monitoring algorithms to compute robust satisfaction values [10, 7, 5], there is limited work on an online monitoring algorithm to do the same [3].

The main technical and theoretical challenge of online monitoring lies in the definition of a practical semantics for a temporal logic formula over a partial signal, i.e., a signal trace with incomplete data which cannot yet validate or invalidate φ . Past work [8] has identified three views for the satisfaction of a LTL property φ over a partial trace τ : (1) a *weak view* where the truth value of φ over τ is assigned to *true* if there is some suffix of τ that satisfies φ , (2) a *strong view* when it is defined to be *false* when some suffix of τ does not satisfy φ and (3) a *neutral view* when the truth value is defined using a truncated semantics of LTL restricted to *finite* paths. In [11], the authors extend the truncated semantics to MTL, and in [3], the authors introduce the notion of a *predictor*, which works as an oracle to complete the partial trace and provide an estimated satisfaction value. However, such a value cannot be formally trusted in general as long as the data is incomplete.

We now outline our major contributions in this paper. In Section 3, we present *robust interval* semantics for an STL property φ on a partial trace τ that unifies the different semantic views of real-time logics on truncated paths. Informally, the robust interval semantics map a trace $\mathbf{x}(t)$ and an STL property φ to an interval (ℓ, v) , with the interpretation that for any suffix $u(t)$, ℓ is the greatest lower bound on the quantitative semantics of the trace $\mathbf{x}(t)$, and v is the corresponding lowest upper bound. There is a natural correspondence between the interval semantics and three-valued semantics: (1) the truth value of φ is false according to the weak view iff v is negative, and true otherwise; (2) the truth value is true according to the strong view iff ℓ is positive, and false otherwise; and (3) a neutral semantics, e.g., based on some predictor, can be defined when $\ell < 0 < v$, i.e., when there exist both suffixes that can violate or satisfy φ .

In Section 4, we present an efficient online algorithm to compute the robust interval semantics for bounded horizon formulas. Our approach is based on the offline algorithm of [5] extended to work in a fashion similar to the incremental Boolean monitoring of STL implemented in the tool AMT [18]. A key feature of our algorithm is that it imposes minimal runtime overhead with respect to the offline algorithm, while being

able to compute robust satisfaction intervals on partial traces. In Section 5, we present specialized algorithms to deal with commonly-used unbounded horizon formulas using only a bounded amount of memory, and give an instance of a unbounded horizon formula for which no robust online monitor with bounded memory exists.

Finally, we present an implementation and experimental results on two large-scale case studies: (i) industrial-scale Simulink models from the automotive domain in Section 6, and (ii) an automatic grading system used in a massive online education initiative on CPS [14]. Since the online algorithm can abort simulation as soon as the truth value of the property is determined, we see a consistent 10%-20% savings in simulation time (which is typically several hours) in a majority of experiments, with negligible overhead (< 1%). In general, our results indicate that the benefits of our online monitoring algorithm over the offline approach far outweigh any overheads.

2 Background

Interval Arithmetic. We make extensive use of simple interval arithmetic operations which we review next. An interval I is a convex subset of \mathbb{R} . A singular interval $[a, a]$ contains exactly one point. Intervals (a, a) , $[a, a)$, $(a, a]$, and \emptyset denote empty intervals. We enumerate interval operations below assuming open intervals. Similar operations can be defined for closed, open-closed, and closed-open intervals. In what follows, let $I_j = (a_j, b_j)$.

$$\begin{aligned} -I_1 &= (-b_1, -a_1) & \min_j(I_j) &= (\min_j(a_j), \min_j(b_j)) \\ c + I_1 &= (c + a_1, c + b_1) & \max_j(I_j) &= (\max_j(a_j), \max_j(b_j)) \\ I_1 \oplus I_2 &= (a_1 + a_2, b_1 + b_2) \\ I_1 \cap I_2 &= \begin{cases} \emptyset & \text{if } \min(b_1, b_2) < \max(a_1, a_2) \\ (\max(a_1, a_2), \min(b_1, b_2)) & \text{otherwise.} \end{cases} \end{aligned} \quad (2.1)$$

Definition 1 (Signal). A time domain \mathcal{T} is a finite or infinite set of time instants such that $\mathcal{T} \subseteq \mathbb{R}^{\geq 0}$ with $0 \in \mathcal{T}$. A signal \mathbf{x} is a function from \mathcal{T} to \mathcal{X} . Given a time domain \mathcal{T} , a partial signal is any signal defined on a time domain $\mathcal{T}' \subseteq \mathcal{T}$.

Simulation frameworks typically provide signal values at discrete time instants, usually this is a by-product of using a numerical technique to solve the differential equations in the underlying system. These discrete-time solutions are assumed to be sampled versions of the actual signal, which can be reconstructed using some form of interpolation. In this paper, we assume constant interpolation to reconstruct the signal $\mathbf{x}(t)$, i.e., given a sequence of time-value pairs $(t_0, \mathbf{x}_0), \dots, (t_n, \mathbf{x}_n)$, for all $t \in [t_0, t_n)$, we define $\mathbf{x}(t) = \mathbf{x}_i$ if $t \in [t_i, t_{i+1})$, and $\mathbf{x}(t_n) = \mathbf{x}_n$. Further, let $\mathcal{T}_n \subseteq \mathcal{T}$ represent the finite subset of time instants at which the signal values are given.

Signal Temporal Logic. We use Signal Temporal Logic (STL) [7] to analyze time-varying behaviors of signals. We now present its syntax and semantics. A *signal predicate* μ is a formula of the form $f(\mathbf{x}) > 0$, where \mathbf{x} is a variable that takes values from \mathcal{X} , and f is a function from \mathcal{X} to \mathbb{R} . For a given f , let f_{\inf} denote $\inf_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$, i.e., the *greatest lower bound* of f over \mathcal{X} . Similarly, let $f_{\sup} = \sup_{\mathbf{x} \in \mathcal{X}} f(\mathbf{x})$. The syntax of an STL formula φ is defined as follows:

$$\varphi ::= \mu \mid \neg\varphi \mid \varphi \wedge \varphi \mid \square_{(u,v)}\varphi \mid \diamond_{(u,v)}\varphi \mid \varphi \mathbf{U}_{(u,v)}\varphi \quad (2.2)$$

Quantitative semantics for timed-temporal logics have been proposed for STL in [7]; we include the definition below.

Definition 2 (Robust Satisfaction Value). *The robust satisfaction value is a function ρ mapping φ , the signal \mathbf{x} , and a time $\tau \in \mathcal{T}$ as follows:*

$$\begin{aligned}
\rho(f(\mathbf{x}) > 0, \mathbf{x}, \tau) &= f(\mathbf{x}(\tau)) \\
\rho(\neg\varphi, \mathbf{x}, \tau) &= -\rho(\varphi, \mathbf{x}, \tau) \\
\rho(\varphi_1 \wedge \varphi_2, \mathbf{x}, \tau) &= \min(\rho(\varphi_1, \mathbf{x}, \tau), \rho(\varphi_2, \mathbf{x}, \tau)) \\
\rho(\Box_I\varphi, \mathbf{x}, \tau) &= \inf_{\tau' \in \tau+I} \rho(\varphi, \mathbf{x}, \tau') \\
\rho(\Diamond_I\varphi, \mathbf{x}, \tau) &= \sup_{\tau' \in \tau+I} \rho(\varphi, \mathbf{x}, \tau') \\
\rho(\varphi \mathbf{U}_I \psi, \mathbf{x}, \tau) &= \sup_{\tau_1 \in \tau+I} \min\left(\rho(\psi, \mathbf{x}, \tau_1), \inf_{\tau_2 \in (\tau, \tau_1)} \rho(\varphi, \mathbf{x}, \tau_2)\right)
\end{aligned} \tag{2.3}$$

Here, the translation from quantitative semantics to the usual Boolean satisfaction semantics is that a signal \mathbf{x} satisfies an STL formula φ at a time τ iff the robust satisfaction value $\rho(\varphi, \mathbf{x}, \tau) \geq 0$.

3 Robust Interval Semantics

In what follows, we assume that we wish to monitor the robust satisfaction value of a signal over a finite time-horizon T_H . We assume that the signal is obtained by applying piecewise constant interpolation to a sampled signal defined over time-instants $\{t_0, t_1, \dots, t_N\}$, such that $t_N = T_H$. In an online monitoring context, at any time t_i , only the partial signal over time instants $\{t_0, \dots, t_i\}$ is available, and the rest of the signal becomes available in discrete time increments. We define robust satisfaction semantics of STL formulas over such partial signals using an interval-based semantics. Such a *robust satisfaction interval* (RoSI) includes all possible robust satisfaction values corresponding to the suffixes of the partial signal. In this section, we formalize the recursive definitions for the robust satisfaction interval of an STL formula with respect to a partial signal, and in the next section we will discuss an efficient algorithm to compute and maintain these intervals.

Definition 3 (Prefix, Completions). *Let $\{t_0, \dots, t_i\}$ be a finite set of time instants such that $t_i \leq T_H$, and let $\mathbf{x}_{[0, i]}$ be a partial signal over the time domain $[t_0, t_i]$. We say that $\mathbf{x}_{[0, i]}$ is a prefix of a signal \mathbf{x} if for all $t \leq t_i$, $\mathbf{x}(t) = \mathbf{x}_{[0, i]}(t)$. The set of completions of a partial signal $\mathbf{x}_{[0, i]}$ (denoted by $\mathcal{C}(\mathbf{x}_{[0, i]})$) is defined as the set $\{\mathbf{x} \mid \mathbf{x}_{[0, i]} \text{ is a prefix of } \mathbf{x}\}$.*

Definition 4 (Robust Satisfaction Interval (RoSI)). *The robust satisfaction interval of an STL formula φ on a partial signal $\mathbf{x}_{[0, i]}$ at a time $\tau \in [t_0, t_N]$ is an interval I such that:*

$$\inf(I) = \inf_{\mathbf{x} \in \mathcal{C}(\mathbf{x}_{[0, i]})} \rho(\varphi, \mathbf{x}, \tau) \quad \text{and} \quad \sup(I) = \sup_{\mathbf{x} \in \mathcal{C}(\mathbf{x}_{[0, i]})} \rho(\varphi, \mathbf{x}, \tau)$$

Definition 5. *We now define a recursive function $[\rho]$ that maps a given formula φ , a partial signal $\mathbf{x}_{[0, i]}$ and a time $\tau \in \mathcal{T}$ to an interval $[\rho](\varphi, \mathbf{x}_{[0, i]}, \tau)$.*

$$\begin{aligned}
[\rho](f(\mathbf{x}_{[0,i]} > 0, \mathbf{x}_{[0,i]}, \tau) &= \begin{cases} [f(\mathbf{x}_{[0,i]}(\tau)), f(\mathbf{x}_{[0,i]}(\tau))] & \tau \in [t_0, t_i] \\ [f_{\text{inf}}, f_{\text{sup}}] & \text{otherwise.} \end{cases} \\
[\rho](\neg\varphi, \mathbf{x}_{[0,i]}, \tau) &= -[\rho](\varphi, \mathbf{x}_{[0,i]}, \tau) \\
[\rho](\varphi_1 \wedge \varphi_2, \mathbf{x}_{[0,i]}, \tau) &= \min([\rho](\varphi_1, \mathbf{x}_{[0,i]}, \tau), [\rho](\varphi_2, \mathbf{x}_{[0,i]}, \tau)) \\
[\rho](\Box_I \varphi, \mathbf{x}_{[0,i]}, \tau) &= \inf_{t \in \tau+I} ([\rho](\varphi, \mathbf{x}_{[0,i]}, t)) \\
[\rho](\Diamond_I \varphi, \mathbf{x}_{[0,i]}, \tau) &= \sup_{t \in \tau+I} ([\rho](\varphi, \mathbf{x}_{[0,i]}, t)) \\
[\rho](\varphi_1 \mathbf{U}_I \varphi_2, \mathbf{x}_{[0,i]}, \tau) &= \sup_{\tau_2 \in \tau+I} \min \left([\rho](\varphi_2, \mathbf{x}_{[0,i]}, \tau_2), \inf_{\tau_1 \in (\tau, \tau_2)} [\rho](\varphi_1, \mathbf{x}_{[0,i]}, \tau_1) \right)
\end{aligned} \tag{3.1}$$

The following lemma that can be proved by induction over the structure of STL formulas shows that the interval obtained by applying the recursive definition for $[\rho]$ is indeed the robust satisfaction interval as defined in Def. 4.

Lemma 1. *For any STL formula φ , the function $[\rho](\varphi, \mathbf{x}_{[0,i]}, \tau)$ defines the robust satisfaction interval for the formula φ over the signal $\mathbf{x}_{[0,i]}$ at time τ .*

4 Online Algorithm

Donzé et al. [5] present an offline algorithm for monitoring STL formulas over (piecewise) linearly interpolated signals. A naïve implementation of an online algorithm is as follows: at time t_i , use a modification of the offline monitoring algorithm to recursively compute the robust satisfaction intervals as defined by Def. 5 to the signal $\mathbf{x}_{[0,i]}$. We observe that such a procedure does many repeated computations that can be avoided by maintaining the results of intermediate computations. Furthermore, the naïve procedure requires storing the signal values over the entire time horizon, which makes it memory-intensive. In this section, we present the main technical contribution of this paper: *an online algorithm that is memory-efficient and avoids repeated computations.*

As in the offline monitoring algorithm in [5], an essential ingredient of the online algorithm is Lemire’s running maximum filter algorithm [16]. The problem this algorithm addresses is the following: given a sequence of values a_1, \dots, a_n , find the maximum (resp. minimum) over all windows of size w , i.e., for all j , $\max_{i \in [j, j+w)} a_i$ (resp. $\min_{i \in [j, j+w)} a_i$). We briefly review an extension of Lemire’s algorithm over piecewise-constant signals with variable time steps, given as Algorithm 1. The main observation in Lemire’s algorithm is that it is sufficient to maintain a descending (resp. ascending) monotonic edge (noted F in Algorithm 1) to compute the sliding maxima (resp. minima), in order to achieve an optimal procedure (measured in terms of the number of comparisons between elements).

We first focus on the fragment of STL where each temporal operator is bounded by a time-interval I such that $\text{sup}(I)$ is finite. The procedure for online monitoring is an algorithm that maintains in memory the syntax tree of the formula φ to be monitored, augmented with some book-keeping information. First, we formalize some notation. For a given formula φ , let \mathcal{T}_φ represent the syntax tree of φ , and let $\text{root}(\mathcal{T}_\varphi)$ denote the root of the tree. Each node in the syntax tree (other than a leaf node) corresponds to

Algorithm 1: SlidingMax($(t_0, \mathbf{x}_0), \dots, (t_N, \mathbf{x}_N), [a, b]$).

Output: Sliding maximum $\mathbf{y}(t)$ over times in $[t_0, t_N]$

```

1  $F := \{0\}$  //  $F$  is the set of times representing the monotonic edge
2  $i := 0; s, t := t_0 - b$ 
3 while  $t + a < t_N$  do
4   if  $F \neq \emptyset$  then  $t := \min(t_{\min(F)} - a, t_{i+1} - b)$ 
5   else  $t := t_{i+1} - b$ 
6   if  $t = t_{i+1} - b$  then
7     while  $\mathbf{x}_{i+1} \geq \mathbf{x}_{\max(F)} \wedge F \neq \emptyset$  do
8        $F := F - \max(F)$ 
9        $F := F \cup \{i + 1\}, i := i + 1$ 
10  else // Slide window to the right
11    if  $s > t_0$  then  $\mathbf{y}(s) := \mathbf{x}_{\min(F)}$ 
12    else  $\mathbf{y}(t_0) := \mathbf{x}_{\min(F)}$ 
13     $F := F - \min(F), s := t$ 

```




Fig. 1. Syntax tree \mathcal{T}_φ for φ (given in (4.2)) with each node v annotated with $\text{hor}(v)$.

an STL operator $\neg, \vee, \wedge, \square_I$ or \diamond_I .³ We will use \mathbf{H}_I to denote any temporal operator bounded by interval I . For a given node v , let $\text{op}(v)$ denote the operator for that node. For any node v in \mathcal{T}_φ (except the root node), let $\text{parent}(v)$ denote the unique parent of v .

Algorithm 2 does the online RoSI computation. Like the offline algorithm, it is a dynamic programming algorithm operating on the syntax tree of the given STL formula, i.e., computation of the RoSI of a formula combines the RoSIs for its constituent subformulas in a bottom-up fashion. As computing the RoSI at a node v requires the RoSIs at the child-nodes, this computation has to be delayed till the RoSIs at the children of v in a certain time-interval are available. We call this time-interval the *time horizon* of v (denoted $\text{hor}(v)$), and define it recursively in Eq. (4.1).

$$\text{hor}(v) = \begin{cases} [0] & \text{if } v = \text{root}(\mathcal{T}_\varphi) \\ I \oplus \text{hor}(\text{parent}(v)) & \text{if } v \neq \text{root}(\mathcal{T}_\varphi) \text{ and } \text{op}(\text{parent}(v)) = \mathbf{H}_I \\ \text{hor}(\text{parent}(v)) & \text{otherwise.} \end{cases} \quad (4.1)$$

We illustrate the working of the algorithm using a small example then give a brief sketch of the various steps in the algorithm.

Example 1. Consider formula (4.2). We show \mathcal{T}_φ and $\text{hor}(v)$ for each node v in \mathcal{T}_φ in Fig. 1. In rest of the paper, we use φ as a running example⁴.

$$\varphi \triangleq \square_{[0,a]} (\neg(y > 0) \vee \diamond_{[b,c]}(x > 0)) \quad (4.2)$$

³ We omit the case of \mathbf{U}_I here for lack of space, although the rewriting approach of [5] can also be adapted and was implemented in our tool.

⁴ We remark that φ is equivalent to $\square_{[0,a]} ((y > 0) \implies \diamond_{[b,c]}(x > 0))$, which is a common formula used to express a timed causal relation between two signals.

The algorithm augments each node v of \mathcal{T}_φ with a double-ended queue, that we denote $\text{worklist}[v]$. Let ψ be the subformula denoted by the tree rooted at v . For the partial signal $\mathbf{x}_{[0,i]}$, the algorithm maintains in $\text{worklist}[v]$, the RoSI $[\rho](\psi, \mathbf{x}_{[0,i]}, t)$ for each $t \in \text{hor}(v) \cap [t_0, t_i]$. We denote by $\text{worklist}[v](t)$ the entry corresponding to time t in $\text{worklist}[v]$. When a new data-point \mathbf{x}_{i+1} corresponding to the time t_{i+1} is available, the monitoring procedure updates each $[\rho](\psi, \mathbf{x}_{[0,i]}, t)$ in $\text{worklist}[v]$ to $[\rho](\psi, \mathbf{x}_{[0,i+1]}, t)$.

In Fig. 3, we give an example of a run of the algorithm. We assume that the algorithm starts in a state where it has processed the partial signal $\mathbf{x}_{[0,2]}$, and show the effect of receiving data at time-points t_3, t_4 and t_5 . The figure shows the states of the worklists at each node of \mathcal{T}_φ at these times when monitoring the STL formula φ presented in Eq. (4.2). Each row in the table adjacent to a node shows the state of the worklist after the algorithm processes the value at the time indicated in the first column.

The first row of the table shows the snapshot of the worklists at time t_2 . Observe that in the worklists for the subformula $y > 0, \neg y > 0$, because $a < b$, the data required to compute the RoSI at t_0, t_1 and the time a , is available, and hence each of the RoSIs is singular. On the other hand, for the subformula $x > 0$, the time horizon is $[b, a + c]$, and no signal value is available at any time in this interval. Thus, at time t_2 , all elements of $\text{worklist}[v_{x>0}]$ are $(\mathbf{x}_{\text{inf}}, \mathbf{x}_{\text{sup}})$ corresponding to the greatest lower bound and lowest upper bound on x .

To compute the values of $\diamond_{[b,c]}(x > 0)$ at any time t , we take the supremum over values from times $t + b$ to $t + c$. As the time horizon for the node corresponding to $\diamond_{[b,c]}(x > 0)$ is $[0, a]$, t ranges over $[0, a]$. In other words, we wish to perform the sliding maximum over the interval $[0 + b, a + c]$, with a window of length $c - b$. We can use the algorithm for computing the sliding window maximum as discussed earlier in this section. One caveat is that we need to store separate monotonic edges for the upper and lower bounds of the RoSIs. The algorithm then proceeds upward on the syntax tree, only updating the worklist of a node only when there is an update to the worklists of its children.

The second row in each table is the effect of obtaining a new time point (at time t_3) for both signals. Note that this does not affect $\text{worklist}[v_{y>0}]$ or $\text{worklist}[v_{\neg y>0}]$, as all RoSIs are already singular, but does update the RoSI values for the node $v_{x>0}$. The algorithm then invokes Alg. 1 on $\text{worklist}[v_{x>0}]$ to update $\text{worklist}[v_{\diamond_{[b,c]}(x>0)}]$. Note that in the invocation on the second row (corresponding to time t_3), there is an additional value in the worklist, at time t_3 . This leads Alg. 1 to produce a new value of $\text{SlidingMax}(\text{worklist}[v_{x>0}], [b, c]) (t_3 - b)$, which is then inserted in $\text{worklist}[v_{\diamond_{[b,c]}(x>0)}]$. This leads to additional points appearing in worklists at the ancestors of this node.

Finally, we remark that the run of this algorithm shows that at time t_4 , the RoSI for the formula φ is $[-2, -2]$, which yields a negative upper bound, showing that the formula is not satisfied irrespective of the suffixes of x and y . In other words, the satisfaction of φ is known before we have all the data required by $\text{hor}(\varphi)$.

Alg. 2 is essentially a procedure that recursively visits each node in the syntax tree \mathcal{T}_φ of the STL formula φ that we wish to monitor. Line 4 corresponds to the base case of the recursion, i.e. when the algorithm visits a leaf of \mathcal{T}_φ or an atomic predicates of the form $f(\mathbf{x}) > 0$. Here, the algorithm inserts the pair $(t_{i+1}, \mathbf{x}_{i+1})$ in $\text{worklist}[v_{f(\mathbf{x})>0}]$ if t_{i+1} lies inside $\text{hor}(v_{f(\mathbf{x})>0})$. In other words, it only tracks a value if it is useful for the computing the robust satisfaction interval of some ancestor node.

For a node corresponding to a Boolean operation, the algorithm first updates the worklists at the children, and then uses them to update the worklist at the node. If




Fig. 2. These plots show the signals $x(t)$ and $y(t)$. Each signal begins at time $t_0 = 0$, and we consider three partial signals: $\mathbf{x}_{[0,3]}$ (black + blue), and $\mathbf{x}_{[0,4]}$ ($\mathbf{x}_{[0,3]}$ + green), and $\mathbf{x}_{[0,5]}$ ($\mathbf{x}_{[0,4]}$ + red).

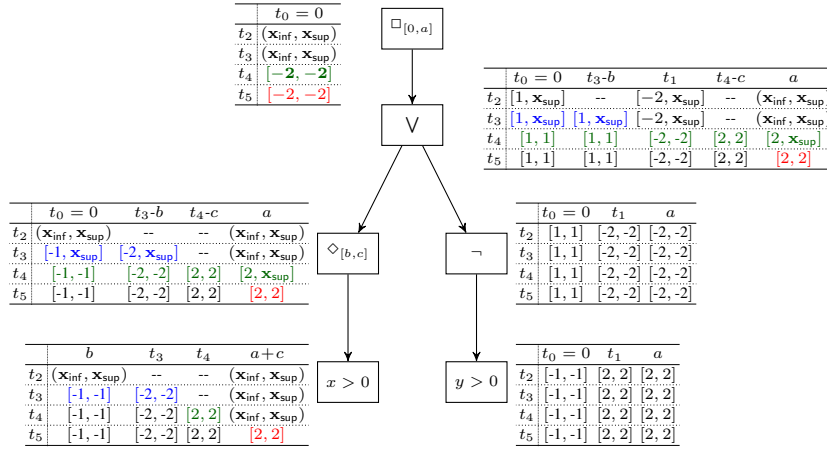


Fig. 3. We show a snapshot of the worklist v maintained by the algorithm for four different (incremental) partial traces of the signals $x(t)$ and $y(t)$. Each row indicates the state of worklist v at the time indicated in the first column. An entry marked -- indicates that the corresponding element did not exist in worklist v at that time. Each colored entry indicates that the entry was affected by availability of a signal fragment of the corresponding color.

the current node represents $\neg\varphi$ (Line 5), the algorithm flips the sign of each entry in worklist v_φ ; this operation is denoted as $\neg\text{worklist}[v_\varphi]$. Consider the case where the current node v_ψ is a conjunction $\varphi_1 \wedge \varphi_2$. The sequence of upper bounds and the sequence of lower bounds of the entries in worklist v_{φ_1} and worklist v_{φ_2} can be each thought of as a piecewise-constant signal (likewise for worklist v_{φ_2}). In Line 11, the algorithm computes a pointwise-minimum over piecewise-constant signals representing the upper and lower bounds of the RoSIs of its arguments. Note that if for $i = 1, 2$, if worklist v_{φ_i} has N_i entries, then the pointwise-min would have to be performed at most $N_1 + N_2$ distinct time-points. Thus, worklist $v_{\varphi_1 \wedge \varphi_2}$ has at most $N_1 + N_2$ entries. A similar phenomenon can be seen in Fig. 3, where computing a max over the worklists of $v_{\diamond_{[b,c]}(x>0)}$ and $v_{\neg(y>0)}$ leads to an increase in the number of entries in the worklist of the disjunction.

For nodes corresponding to temporal operators, e.g., $\diamond_I\varphi$, the algorithm first updates worklist v_φ . It then applies Alg. 1 to compute the sliding maximum over worklist v_φ . Note that if worklist v_φ contains N entries, so does worklist $v_{\diamond_I\varphi}$.

Algorithm 2: $\text{updateWorkList}(v_\psi, t_{i+1}, \mathbf{x}_{i+1})$

```
//  $v_\psi$  is a node in the syntax tree,  $(t_{i+1}, \mathbf{x}_{i+1})$  is a new
signal time-point
1 switch  $\psi$  do
2   case  $f(\mathbf{x}) > 0$  do
3     if  $t_{i+1} \in \text{hor}(v_\psi)$  then
4        $\text{worklist}[v_\psi](t_{i+1}) := [f(\mathbf{x}_{i+1}), f(\mathbf{x}_{i+1})]$ 
5   case  $\neg\varphi$  do
6      $\text{updateWorkList}(v_\varphi, t_{i+1}, \mathbf{x}_{i+1})$ ;
7      $\text{worklist}[v_\psi] := -\text{worklist}[v_\varphi]$ 
8   case  $\varphi_1 \wedge \varphi_2$  do
9      $\text{updateWorkList}(v_{\varphi_1}, t_{i+1}, \mathbf{x}_{i+1})$ ;
10     $\text{updateWorkList}(v_{\varphi_2}, t_{i+1}, \mathbf{x}_{i+1})$ ;
11     $\text{worklist}[v_\psi] := \min(\text{worklist}[v_{\varphi_1}], \text{worklist}[v_{\varphi_2}])$ 
12   case  $\Box_I \varphi$  do
13      $\text{updateWorkList}(v_\varphi, t_{i+1}, \mathbf{x}_{i+1})$ ;
14      $\text{worklist}[v_\psi] := \text{SlidingMax}(\text{worklist}[v_\varphi], I)$ 
```

A further optimization can be implemented on top of this basic scheme. For a node v corresponding to the subformula $\mathbf{H}_I \varphi$, the first few entries of $\text{worklist}[v]$ (say up to time u) could become singular intervals once the required RoSIs for $\text{worklist}[v_\varphi]$ are available. The optimization is to only compute SlidingMax over $\text{worklist}[v_\varphi]$ starting from $u + \text{inf}(I)$. We omit the pseudo-code for brevity.

5 Monitoring untimed formulas

If the STL formula being monitored has untimed (i.e. infinite-horizon) temporal operators, a direct application of Algorithm 2 requires every node in the sub-tree rooted at the untimed operator to have an unbounded time horizon. In other words, for all such nodes, the algorithm would have to keep track of every value over arbitrarily long intervals. For certain untimed operators and the combinations thereof, we show that we can monitor the formulas using only a bounded amount of information. Let $\text{sub}(\varphi)$ denote the set of all subformulas of φ except φ itself. Let $\text{last}(\varphi)$ be defined as follows:

$$\text{last}(\varphi) \triangleq \max_{\psi \in \text{sub}(\varphi)} \sup(\text{hor}(\psi)) \quad (5.1)$$

The meaning of $\text{last}(\varphi)$ is as follows: the last time at which a data value of \mathbf{x} is required to compute $\rho(\varphi, \mathbf{x}, t)$, is $t + \text{last}(\varphi)$. For the formula φ defined in Eq. (4.2), $\text{last}(\varphi) = a + c$. For the formula $\psi \equiv \Box(x > 0)$, $\text{last}(\psi) = \infty$. In general, for any untimed formula φ , $\text{last}(\varphi)$ is equal to ∞ . In Theorem 1, we show that certain classes of untimed formulas can be monitored in an online fashion with bounded amount of memory. We first define the following quantities:

$$\Delta \triangleq \min_{i \geq 0} (t_{i+1} - t_i) \quad w_\varphi \triangleq \max_{\psi \in \text{sub}(\varphi)} \text{last}(\psi) \quad k_\varphi \triangleq \lceil \frac{w_\varphi}{\Delta} \rceil. \quad (5.2)$$

Here, Δ represents the smallest time-step in the monitored signal, w_φ is the largest time horizon of all subformulas of φ , and k_φ is the largest number of discrete time-points for the trace in any w_φ interval.




Fig. 4. A depiction of the action of the procedure to update the summary while computing $[\rho](\square\psi, \mathbf{x}_{[0,i]}, t_0)$. Here, $W(j)$ is shorthand for $[\rho](\psi, \mathbf{x}_{[0,i]}, t_j)$ and $W'(j)$ is shorthand for $[\rho](\psi, \mathbf{x}_{[0,i+1]}, t_j)$.

Theorem 1. *If w_φ is finite, then for each φ listed below, we can monitor RoSI of φ in an online fashion using $O(k_\varphi)$ memory.*

1. $\square\psi(\text{dually } \diamond\psi)$
 2. $\varphi\mathbf{U}\psi$
 3. $\square\diamond\psi(\text{dually } \diamond\square\psi)$
 4. $\square(\varphi \vee \diamond\psi)(\text{dually } \diamond(\varphi \wedge \square\psi))$,
- (5.3)

Proof. Due to space constraints, we only provide proof sketches. The main argument in each of the proofs is as follows: For any partial signal $\mathbf{x}_{[0,i]}$, there are two cases: The first case is when $t_0 \geq t_i - w_\varphi$. By assumption, there are at most k_φ time-points in the interval $[t_0, t_i]$. Thus, in this case, the worklists at each of the nodes v_ψ corresponding to $\psi \in \text{sub}(\varphi)$ have to track at most k_φ RoSI values in order to compute $[\rho](\varphi, \mathbf{x}, t_0)$.

The second case is when $t_0 < t_i - w_\varphi$; this implies that there is a largest time t_{s_i} in $[t_0, t_1, \dots, t_i]$ such that $t_{s_i} < t_i - w_\varphi$. For the partial signal $\mathbf{x}_{[0,i]}$, at each time $t \leq t_{s_i}$, there is enough information to compute the exact robustness value of each of the subformulas of φ . The central step is that for each of the formulas mentioned above, the robustness values in the interval $[t_0, t_{s_i}]$ can be *summarized* to a single robustness value. Furthermore, the interval $(t_{s_i}, t_i]$ can have at most k_φ time-points. Thus, the computation of $[\rho](\varphi, \mathbf{x}_{[0,i]}, t_0)$ can be split into tracking a summary for the interval $[t_{s_i}, t_i]$ and tracking at most k_φ RoSIs in the worklists of the immediate subformulas of φ in the interval $(t_{s_i}, t_i]$. We now explain how the summary information is maintained for each formula.

(1) $[\square\psi]$ We maintain the summary $S = \inf_{j \in [0, s_i]} [\rho](\psi, \mathbf{x}_{[0,i]}, t_j)$, i.e., the infimum over all exact robustness values computable over the partial signal $\mathbf{x}_{[0,i]}$. When a new time-point $(t_{i+1}, \mathbf{x}_{i+1})$ becomes available, S is updated if there is a new time $t_{s_{i+1}}$ for which $[\rho](\psi, \mathbf{x}_{[0,i]}, t_{s_{i+1}})$ can be exactly computed; otherwise, the new value is used to update all entries $[\rho](\psi, \mathbf{x}_{[0,i+1]}, t_j)$ for $t_j \in [t_{s_{i+1}}, t_i]$, and a new entry corresponding to time t_{i+1} is added to $\text{worklist}[v_\psi]$. Please see Fig. 4 for a depiction of this step. We then establish the following: (1) There are at most k_φ entries (each corresponding to $[\rho](\psi, \mathbf{x}_{[0,i]}, t_j)$ for $t_j \in (t_{s_i}, t_i]$) in $\text{worklist}[v_\psi]$. This is true because there can be at most k_φ consecutive time-points that do not update S in any interval of length w_φ . (2) We show by induction that the inf of S and the k_φ entries in $\text{worklist}[v_\psi]$ is equal to $[\rho](\square\psi, \mathbf{x}_{[0,i]}, t_0)$.

Algorithm 3: Computing RoSI for untimed Until

```
1  $v_1 := S_\varphi, v_2 := S_\psi$ 
2 foreach  $j \in [s_i + 1, i]$  do
3    $v_1 := \min(v_1, [\rho](\varphi, \mathbf{x}_{[0,i]}, t_j))$ 
4    $v_2 := \sup(v_2, \min(v_1, [\rho](\psi, \mathbf{x}_{[0,i]}, t_j)))$ 
5  $[\rho](\varphi \mathbf{U} \psi, \mathbf{x}_{[0,i]}, t_0) := v_2$ 
```

Algorithm 4: Computing RoSI for $\square(\varphi \vee \diamond\psi)$

```
1  $v := S$ 
2 foreach  $j \in [s_i + 1, i]$  do
3    $v := \sup([\rho](\psi, \mathbf{x}_{[0,i]}, t_j), \inf(v, [\rho](\varphi, \mathbf{x}_{[0,i]}, t_j)))$ 
4  $[\rho](\square(\varphi \vee \diamond\psi), \mathbf{x}_{[0,i]}, t_0) := v$ 
```

(2) $[\varphi \mathbf{U} \psi]$ We maintain the following two quantities as the summary:

(a) $S_\varphi = [\rho](\square_{[0, t_{s_i}]} \varphi, \mathbf{x}_{[0,i]}, t_0)$ and (b) $S_\psi = [\rho](\varphi \mathbf{U}_{[0, t_{s_i}]} \psi, \mathbf{x}_{[0,i]}, t_0)$. In $\text{worklist}[v_\varphi]$ and $\text{worklist}[v_\psi]$ we store at most k_φ values corresponding to $[\rho](\varphi, \mathbf{x}_{[0,i]}, t_j)$ and $[\rho](\psi, \mathbf{x}_{[0,i]}, t_j)$ for $t_j \in (t_{s_i}, t_i]$. The crucial step is to combine S_φ and S_ψ with the entries in $\text{worklist}[v_\varphi]$ and $\text{worklist}[v_\psi]$ to obtain $[\rho](\varphi \mathbf{U} \psi, \mathbf{x}_{[0,i]}, t_0)$. We show that the iterative procedure in Algorithm 3 can accomplish this. In its j^{th} iteration v_1 is equal to $\inf_{\ell \in [0, j]} [\rho](\varphi, \mathbf{x}_{[0,i]}, t_\ell)$,

and we can show by induction that v_2 is equal to $\sup_{m \in [0, j]} \min \left([\rho](\psi, \mathbf{x}_{[0,i]}, t_m), \inf_{\ell \in [0, m]} [\rho](\varphi, \mathbf{x}_{[0,i]}, t_\ell) \right)$.

Thus, at the end of the computation, the value computed in v_2 is $[\rho](\varphi \mathbf{U} \psi, \mathbf{x}_{[0,i]}, t_0)$.

(3) $[\square \diamond \psi]$ We show that we do not need any additional storage for monitoring φ . Concretely, we posit that $[\rho](\square \diamond \psi, \mathbf{x}_{[0,i]}, t_0) = [\rho](\psi, \mathbf{x}_{[0,i]}, t_i)$. We successively rewrite $[\rho](\square \diamond \psi, \mathbf{x}_{[0,i]}, t_0) = \inf_{j \in [0, i]} \sup_{\ell \in [j, i]} [\rho](\psi, \mathbf{x}_{[0,i]}, t_\ell)$ as follows:

$$\inf \left(\sup_{\ell \in [0, i]} [\rho](\psi, \mathbf{x}_{[0,i]}, t_\ell), \sup_{\ell \in [1, i]} [\rho](\psi, \mathbf{x}_{[0,i]}, t_\ell), \dots, \sup_{\ell \in [i, i]} [\rho](\psi, \mathbf{x}_{[0,i]}, t_\ell) \right) \quad (5.4)$$

$$\inf \left(\sup([\rho](\psi, \mathbf{x}_{[0,i]}, t_i), \sup_{\ell \in [0, i-1]} [\rho](\psi, \mathbf{x}_{[0,i]}, t_\ell)), \sup([\rho](\psi, \mathbf{x}_{[0,i]}, t_i), \sup_{\ell \in [1, i-1]} [\rho](\psi, \mathbf{x}_{[0,i]}, t_\ell)), \dots, [\rho](\psi, \mathbf{x}_{[0,i]}, t_i) \right) \quad (5.5)$$

In the above, to go from (5.4) to (5.5), we expand the inner sup expressions, and observe that the last term in the inf evaluates to $[\rho](\psi, \mathbf{x}_{[0,i]}, t_i)$. For the final step, we observe that $\inf(I_1, \sup(I_1, I_2), \dots, \sup(I_1, I_n)) = I_1$, and thus, (5.5) simplifies to $[\rho](\psi, \mathbf{x}_{[0,i]}, t_i)$. By duality, a similar proof works for $\diamond \square \psi$.

(4) $[\square(\varphi \vee \diamond\psi)]$ We maintain one quantity as the summary information: $S = [\rho](\square_{[0, s_i]}(\varphi \vee \diamond\psi), \mathbf{x}_{[0,i]}, t_0)$. Additionally, we store at most k_φ entries corresponding to $[\rho](\varphi, \mathbf{x}_{[0,i]}, t_j)$ in $\text{worklist}[v_\varphi]$ and at most k_ψ entries corresponding to $[\rho](\psi, \mathbf{x}_{[0,i]}, t_j)$ in $\text{worklist}[v_\psi]$. To compute $[\rho](\varphi, \mathbf{x}_{[0,i]}, t_0)$, we use Algorithm 4. To complete the proof we observe that Algorithm 4 computes expression (5.6) that has nested and alternating sups and infs:

$$\sup([\rho](\psi, \mathbf{x}_{[0,i]}, t_i), \inf([\rho](\varphi, \mathbf{x}_{[0,i]}, t_i), \sup([\rho](\psi, \mathbf{x}_{[0,i]}, t_{i-1}), \dots, S])) \quad (5.6)$$

Using the identity $\sup(I_1, \inf(I_2, I_3)) = \inf(\sup(I_1, I_2), \sup(I_1, I_3))$, we can rearrange the above expression to obtain:

$$\inf \left(\sup \left([\rho](\psi, \mathbf{x}_{[0,i]}, t_i), [\rho](\varphi, \mathbf{x}_{[0,i]}, t_i) \right), \sup \left([\rho](\psi, \mathbf{x}_{[0,i]}, t_i), [\rho](\psi, \mathbf{x}_{[0,i]}, t_{i-1}), \inf \left([\rho](\varphi, \mathbf{x}_{[0,i]}, t_{i-1}), \sup \left([\rho](\psi, \mathbf{x}_{[0,i]}, t_{i-2}), \dots, S \right) \right) \right) \right) \quad (5.7)$$

By repeated use of this identity on the expression in the second line, we get the expres-

sion $\inf_{j \in [0,i]} \left(\max \left([\rho](\varphi, \mathbf{x}_{[0,i]}, t_j), \sup_{\ell \in [j,i]} [\rho](\psi, \mathbf{x}_{[0,i]}, t_\ell) \right) \right)$, which is equal to $[\rho](\varphi, \mathbf{x}_{[0,i]}, t_0)$. ■

Lemma 2. *There is an STL formula φ and a trace $\mathbf{x}_1, \dots, \mathbf{x}_n$, such that computing $[\rho](\varphi, \mathbf{x}_{[0,i]}, t_0)$ for any i requires storing $O(i)$ values.*

Proof. Consider the signal $\mathbf{x} = (x, y)$. Consider the formula: $\varphi \equiv \diamond((x > 0) \wedge \diamond(y > 0))$. By definition, $[\rho](\varphi, \mathbf{x}_{[0,i]}, t_0) =$

$$\sup_{j \in [0,i]} \left(\inf(x_0, y_0, y_1, \dots, y_i), \inf(x_1, y_1, y_2, \dots, y_i), \dots, \inf(x_i, y_i) \right) \quad (5.8)$$

Suppose for all $j \in [0, i-1]$, $x_j > x_{j+1}$, for all $j \in [0, i-1]$, $y_j > y_{j+1}$. Furthermore, suppose that there exists some k , such that for all $j \in [0, k]$, $y_j > x_j$. The two signals satisfying these conditions are shown in Fig. 5. For the j^{th} term inside the outer sup in (5.8), when data at time t_i is available, as long as $x_j < \inf_{\ell \in [j,i]} y_\ell$, we need to store the value of x_j , because there can be a future value y_p such that $y_p < x_j$. Furthermore, we also need to store all values $\inf_{\ell \in [j,i]} y_\ell$ as the new value y_p may be smaller than only some x_j . As shown in Fig. 5, there can be arbitrarily long intervals of length m , where this condition is true; i.e., we need to store at least $O(m)$ values of x and y . For very large values of m , observe that $m \approx i$, which completes the proof. ■




Fig. 5. Two signals demonstrating that memory proportional to the size of the signal is required for monitoring the formula $\diamond(\varphi \wedge \diamond\psi)$

6 Experimental Results

We implemented Algorithm 2 as a stand-alone tool that can be plugged in loop with any black-box simulator and evaluated it using two practical real-world applications. We considered the following criteria: (1) On an average, what fraction of simulation time can be saved by online monitoring? (2) How much overhead does online monitoring add, and how does it compare to a naïve implementation that at each step recomputes everything using an offline algorithm?

6.1 Diesel Engine Model (DEM)

The first case study is an industrial-sized Simulink[®] model of a prototype airpath system in a diesel engine. The closed-loop model consists of a plant model describing the airpath dynamics, and a controller implementing a proprietary control scheme. The model has more than 3000 blocks, with more than 20 lookup tables approximating high-dimensional nonlinear functions. Due to the significant model complexity, the speed of simulation is about 5 times slower, i.e., simulating 1 second of operation takes 5 seconds in Simulink[®]. As it is important to simulate this model over a long time-horizon to characterize the airpath behavior over extended periods of time, savings in simulation-time by early detection of requirement violations is very beneficial. We selected two parameterized safety requirements after discussions with the control designers, (shown in Eq. (6.1)-(6.2)). Due to proprietary concerns, we suppress the actual values of the parameters used in the requirements.

$$\varphi_{overshoot}(\mathbf{p}_1) = \square_{[a,b]}(\mathbf{x} < c) \quad (6.1)$$

$$\varphi_{transient}(\mathbf{p}_2) = \square_{[a,b]}(|\mathbf{x}| > c \implies (\diamond_{[0,d]}|\mathbf{x}| < e)) \quad (6.2)$$

Property $\varphi_{overshoot}$ with parameters $\mathbf{p}_1 = (a, b, c)$ specifies that in the interval $[a, b]$, the overshoot on the signal \mathbf{x} should remain below a certain threshold c . Property $\varphi_{transient}$ with parameters $\mathbf{p}_2 = (a, b, c, d, e)$ is a specification on the settling time of the signal \mathbf{x} . It specifies that in the time interval $[a, b]$ if at some time t , $|\mathbf{x}|$ exceeds c then it settles to a small region ($|\mathbf{x}| < e$) before $t + d$. In Table 1, we consider three different valuations ν_1, ν_2, ν_3 for \mathbf{p}_1 in the requirement $\varphi_{overshoot}(\mathbf{p}_1)$, and two different valuations ν_4, ν_5 for \mathbf{p}_2 in the requirement $\varphi_{transient}(\mathbf{p}_2)$.

The main reason for the better performance of the online algorithm is that simulations are time-consuming for this model. The online algorithm can terminate a simulation earlier (either because it detected a violation or obtained a concrete robust satisfaction interval), thus obtaining significant savings. For $\varphi_{overshoot}(\nu_3)$, we choose the parameter values for a and b such that the online algorithm has to process the entire signal trace, and is thus unable to terminate earlier. Here we see that the total overhead (in terms of runtime) incurred by the extra book-keeping by Algorithm 2 is negligible (about 0.1%).

6.2 CPSGrader

CPSGrader [14,6] is a publicly-available automatic grading and feedback generation tool for online virtual labs in cyber-physical systems. It employs temporal logic based testers to check for common fault patterns in student solutions for lab assignments.

Requirement	Num. Traces	Early Termination	Simulation Time (hours)	
			Offline	Online
$\varphi_{overshoot}(\nu_1)$	1000	801	33.3803	26.1643
$\varphi_{overshoot}(\nu_2)$	1000	239	33.3805	30.5923
$\varphi_{overshoot}(\nu_3)$	1000	0	33.3808	33.4369
$\varphi_{transient}(\nu_4)$	1000	595	33.3822	27.0405
$\varphi_{transient}(\nu_5)$	1000	417	33.3823	30.6134

Table 1. Experimental results on DEM.

CPSGrader uses the National Instruments Robotics Environment Simulator to generate traces from student solutions and monitors STL properties (each corresponding to a particular faulty behavior) on them. In the published version of CPSGrader [14], this is done in an offline fashion by first running the complete simulation until a pre-defined cut-off and then monitoring the STL properties on offline traces. At a step-size of 5 ms, simulating 6 sec. of real-world operation of the system takes 1 sec. for the simulator. When students use CPSGrader for active feedback generation and debugging, simulation constitutes the major chunk of the application response time. Online monitoring helps in reducing the response time by avoiding unnecessary simulations, giving the students feedback as soon as faulty behavior is detected.

We evaluated our online monitoring algorithm, on the traces and STL properties used in the published version of CPSGrader [14, 6]. These traces are the result of running actual student submissions on a battery of tests. For lack of space, we refer the reader to [14] for details about the tests and STL properties. As an illustrative example, we show the `keep_bump` property in Eq. 6.3:

$$\varphi_{\text{keep_bump}} = \diamond_{[0,60]} \square_{[0,5]} (\text{bump_right}(t) \vee \text{bump_left}(t)) \quad (6.3)$$

For each STL property, Table 2 compares the total simulation time needed for both the online and offline approaches, summed over all traces. For the offline approach, a suitable simulation cut-off time of 60 sec. is chosen. At a step-size of 5 ms, each trace is roughly of length 1000. For the online algorithm, simulation terminates before this cut-off if the truth value of the property becomes known, otherwise it terminates at the cut-off. Table 2 also shows the monitoring overhead incurred by a naïve online algorithm that performs complete recomputation at every step against the overhead incurred by Alg. 2. Table 2 demonstrates that online monitoring ends up saving up to 24% simulation time ($> 10\%$ in a majority of cases). The monitoring overhead of Alg. 2 is negligible ($< 1\%$) as compared to the simulation time and it is less than the overhead of the naïve online approach consistently by a factor of 40x to 80x.

References

1. Y. Annpureddy, C. Liu, G. Fainekos, and S. Sankaranarayanan. S-TaLiRo: A tool for temporal logic falsification for hybrid systems. In *TACAS*, pages 254–257. 2011.
2. E. Bartocci, L. Bortolussi, and G. Sanguinetti. Data-driven statistical learning of temporal logic properties. In *Formal Modeling and Analysis of Timed Systems*, pages 23–37. Springer International Publishing, 2014.
3. A. Dokhanchi, B. Hoxha, and G. Fainekos. On-line monitoring for temporal logic robustness. In *RV*, pages 231–246. 2014.

STL Test Bench	Num. Traces	Early Termination	Sim. Time (mins)		Overhead (secs)	
			Offline	Online	Naive	Alg. 2
avoid_front	1776	466	296	258	553	9
avoid_left	1778	471	296	246	1347	30
avoid_right	1778	583	296	226	1355	30
hill_climb ₁	1777	19	395	394	919	11
hill_climb ₂	1556	176	259	238	423	7
hill_climb ₃	1556	124	259	248	397	7
filter	1451	78	242	236	336	6
keep_bump	1775	468	296	240	1.2×10^4	268
what_hill	1556	71	259	253	1.9×10^4	1.5×10^3

Table 2. Evaluation of online monitoring for CPSGrader. Each STL Test Bench has an associated STL property.

4. A. Donzé. Breach, a toolbox for verification and parameter synthesis of hybrid systems. In *CAV*, pages 167–170, 2010.
5. A. Donzé, T. Ferrère, and O. Maler. Efficient robust monitoring for STL. In *CAV*, pages 264–279, 2013.
6. A. Donzé, G. Juniwal, J. C. Jensen, and S. A. Seshia. Cpsgrader website. <http://www.cpsgrader.org>.
7. A. Donzé and O. Maler. Robust satisfaction of temporal logic over real-valued signals. In *Formal modeling and analysis of timed systems*, pages 92–106. 2010.
8. C. Eisner, D. Fisman, J. Havlicek, Y. Lustig, A. McIsaac, and D. V. Campenhout. Reasoning with temporal logic on truncated paths. In *CAV*, pages 27–39, 2003.
9. G. Fainekos, S. Sankaranarayanan, K. Ueda, and H. Yazarel. Verification of automotive control applications using s-taliro. In *Proc. of the American Control Conference*, 2012.
10. G. E. Fainekos and G. J. Pappas. Robustness of temporal logic specifications for continuous-time signals. *Theor. Comp. Sci.*, 410(42):4262–4291, 2009.
11. H.-M. Ho, J. Ouaknine, and J. Worrell. Online monitoring of metric temporal logic. In *Runtime Verification*. 2014.
12. B. Hoxha, H. Abbas, and G. Fainekos. Benchmarks for temporal logic requirements for automotive systems. In *Proc. of Applied Verification for Continuous and Hybrid Systems*, 2014.
13. X. Jin, A. Donzé, J. V. Deshmukh, and S. A. Seshia. Mining requirements from closed-loop control models. In *Proc. of HSCC*, pages 43–52, 2013.
14. G. Juniwal, A. Donzé, J. C. Jensen, and S. A. Seshia. CPSGrader: Synthesizing temporal logic testers for auto-grading an embedded systems laboratory. In *EMSOFT*, October 2014.
15. Z. Kong, A. Jones, A. Medina Ayala, E. Aydin Gol, and C. Belta. Temporal logic inference for classification and prediction from data. In *Proceedings of the 17th international conference on Hybrid systems: computation and control*, pages 273–282. ACM, 2014.
16. D. Lemire. Streaming maximum-minimum filter using no more than three comparisons per element. *arXiv preprint cs/0610046*, 2006.
17. O. Maler and D. Nickovic. Monitoring temporal properties of continuous signals. In *FORMATS/FTRTFT*, pages 152–166, 2004.
18. D. Nickovic and O. Maler. AMT: A property-based monitoring tool for analog systems. *Formal Modeling and Analysis of Timed Systems*, pages 304–319, 2007.