

Model predictive online monitoring of dynamical systems for nested signal temporal logic specifications^{☆,☆☆}

Tao Han, Shaoyuan Li, Xiang Yin^{*}

School of Automation and Intelligent Sensing, Shanghai Jiao Tong University, Shanghai 200240, China

ARTICLE INFO

Keywords:

Signal temporal logic
Syntax tree
Online monitoring
Model prediction

ABSTRACT

This paper investigates the online monitoring problem for cyber–physical systems under signal temporal logic (STL) specifications. The objective is to design an online monitor that evaluates system correctness at runtime based on partial signal observations up to the current time so that alarms can be issued whenever the specification is violated or will inevitably be violated in the future. We consider a model-predictive setting where the system’s dynamic model is available and can be leveraged to enhance monitoring accuracy. However, existing approaches are limited to a restricted class of STL formulae, permitting only a single application of temporal operators. This work addresses the challenge of nested temporal operators in the design of model-predictive monitors. Our method utilizes syntax tree structures to resolve dependencies between temporal operators and introduces the concept of basic satisfaction vectors. A new model-predictive monitoring algorithm is proposed by recursively updating these vectors online while incorporating pre-computed satisfaction regions derived from offline model analysis. We prove that the proposed approach is both sound and complete, ensuring no false or missed alarms. Case studies are provided to demonstrate the effectiveness of our method.

1. Introduction

Specification-based monitoring has emerged as a popular approach for evaluating the safety and correctness of complex engineering cyber–physical systems, such as smart cities [1], autonomous vehicles [2], and industrial IoTs [3]. In this context, a monitor observes the state trajectory generated by the system and evaluates the correctness of the trajectory based on a given formal specification [4–6]. Compared with formal verification techniques such as model checking, the key advantage of monitoring is that it is more lightweight, as it only assesses the correctness of the system along a single observed trajectory without explicitly enumerating the entire reachable state space. Therefore, monitoring can be designed independently and implemented as an add-on module for arbitrary systems—even those treated as black boxes.

In recent years, significant advancements have been made in designing monitoring algorithms for various types of formal specifications, such as Linear Temporal Logic (LTL) [7], Metric Temporal Logic (MTL) [8,9], and Signal Temporal Logic (STL) [10,11]. Among these, STL has become one of the most widely used formal specifications for CPS due to its ability to characterize complex spatio-temporal constraints in real-valued, real-time physical signals. Monitoring algorithms can be further categorized into offline monitoring [9,12] and online monitoring [8,10], depending on the information available to the monitor. In the offline setting, the monitor evaluates the correctness of a system using the entire trajectory. However, this approach is unsuitable for systems operating

[☆] This article is part of a Special issue entitled: ‘TCL3 Discrete-event and hybrid systems(IFAC WC 2026)’ published in Nonlinear Analysis: Hybrid Systems.

^{☆☆} This work was supported by the National Natural Science Foundation of China (62573291,62533017,62173226).

^{*} Corresponding author.

E-mail addresses: coolhantao@sjtu.edu.cn (T. Han), syli@sjtu.edu.cn (S. Li), yinxiang@sjtu.edu.cn (X. Yin).

<https://doi.org/10.1016/j.nahs.2026.101772>

Received 1 November 2025; Received in revised form 6 March 2026; Accepted 29 May 2026

Available online 10 June 2026

1751-570X/© 2026 Elsevier Ltd. All rights are reserved, including those for text and data mining, AI training, and similar technologies.

in real time, where only the signal generated up to the current moment is accessible. In contrast, online monitoring algorithms must account for all possible future behaviors to issue early warnings or terminate processes before a specification violation occurs. As a result, online monitoring is widely adopted in safety-critical systems as a predictive add-on component to ensure operational safety.

Due to the physical dynamics of the system, online signals are not arbitrary in general but follow underlying rules. By leveraging information about system behavior, we can better reason about the feasibility or likelihood of future signals, enabling more precise monitoring decisions [13]. To this end, *predictive monitoring* has recently emerged for systems with either an explicit model or a large dataset of operational history. For example, in [14–17], neural networks trained on historical data are used to predict future signals in order to improve monitoring accuracy. However, since these approaches adopt a data-driven methodology, their predictions may be unreliable. To address this, uncertainty quantification techniques, such as conformal predictions, are often employed to ensure prediction confidence [18–20]. While effective, these methods may introduce additional conservatism into the final monitoring results.

More recently, by leveraging the explicit system model of the underlying dynamical system, a new framework called *model-predictive monitoring* has been proposed in [21]. This approach further improves monitoring accuracy when a precise system model is available. By integrating offline reachability computations with online satisfaction evaluation, it achieves an effective balance between computational complexity and accuracy. In [22], a self-triggered information acquisition mechanism is introduced for model-predictive monitoring to reduce state-sampling overhead. However, existing model-predictive monitoring algorithms remain limited to a simple fragment of STL specifications and nested temporal operators, e.g., “eventually always stays in a region”, are currently unsupported.

In this work, we investigate the design of model-predictive online monitors for signal temporal logic specifications. Unlike existing approaches, we consider a general fragment of STL that permits nested temporal operators. This nested setting introduces fundamental challenges beyond prior works, which track STL formula progress using an index set of remaining formulae. However, this approach fails under nested temporal constraints due to dependencies between operators. To address this, we propose to use syntax trees [23–25] to resolve temporal operator dependencies, and use basic satisfaction vectors as dynamically updated key information during online monitoring. Our algorithm leverages these vectors, updated recursively in real time, along with pre-computed safety regions derived from offline model analysis. This framework successfully extends model-predictive monitoring to general STL specifications. Case studies are also provided to demonstrate the effectiveness of our approach.

2. Preliminary

2.1. System model

We consider a discrete-time control system of form

$$x_{k+1} = f(x_k, u_k), \quad (1)$$

where $x_k \in \mathcal{X} \subset \mathbb{R}^n$ denotes the system state at time instant $k \in \mathbb{Z}_{\geq 0}$, $u_k \in \mathcal{U} \subset \mathbb{R}^m$ represents the control input at time instant k and $f : \mathcal{X} \times \mathcal{U} \rightarrow \mathcal{X}$ is the dynamic function of the system, assumed to be continuous in $\mathcal{X} \times \mathcal{U}$. We assume that the state space \mathcal{X} and input space \mathcal{U} are both bounded due to physical constraints.

Suppose that the system state is x_k at time instant k . Then given a sequence of control inputs $\mathbf{u}_{k:T-1} = u_k u_{k+1} \dots u_{T-1} \in \mathcal{U}^{T-k}$, the corresponding state trajectory generated by the system is defined as $\xi_f(x_k, \mathbf{u}_{k:T-1}) = \mathbf{x}_{k+1:T} = x_{k+1} \dots x_T \in \mathcal{X}^{T-k}$, where each subsequent state satisfies the recursive relation $x_{i+1} = f(x_i, u_i)$ for all $i = k, \dots, T-1$.

2.2. Signal temporal logic

We adopt Signal Temporal Logic (STL) as the formal specification language to evaluate trajectory correctness. The syntax of STL formulae is recursively defined as:

$$\Phi ::= \top \mid \pi^\mu \mid \neg\Phi \mid \Phi_1 \wedge \Phi_2 \mid \Phi_1 \mathbf{U}_{[a,b]} \Phi_2,$$

where \top is the *true* predicate and π^μ is an atomic predicate whose truth value is determined by the sign of its underlying predicate function $\mu : \mathbb{R}^n \rightarrow \mathbb{R}$. Specifically, an atomic predicate π^μ is true at state x_k when $\mu(x_k) \geq 0$; otherwise it is false. Operators \neg and \wedge are the standard Boolean operators “negation” and “conjunction”, respectively. One can further use them to induce other operators such as “disjunction” $\Phi_1 \vee \Phi_2 := \neg(\neg\Phi_1 \wedge \neg\Phi_2)$ and “implication” $\Phi_1 \rightarrow \Phi_2 := \neg\Phi_1 \vee \Phi_2$. $\mathbf{U}_{[a,b]}$ is the temporal operator “until”, where $a, b \in \mathbb{Z}_{\geq 0}$ are two integers with $a \leq b$. Note that, since we consider discrete-time setting, $[a, b]$ is the set of all integers between a and b including themselves.

Let $\mathbf{x} = x_0 x_1 \dots$ be a state sequence, $k \in \mathbb{Z}_{\geq 0}$ be a time instant and Φ be an STL formula. We denote by $(\mathbf{x}, k) \models \Phi$ if sequence \mathbf{x} satisfies STL formula Φ at time instant k . The reader is referred to [26] for more details on the semantics of STL formulae. Particularly, for atomic predicates, we have $(\mathbf{x}, k) \models \pi^\mu$ iff $\mu(x_k) \geq 0$, i.e., $\mu(x_k)$ is non-negative for the current state x_k , and for temporal operators, we have $(\mathbf{x}, k) \models \Phi_1 \mathbf{U}_{[a,b]} \Phi_2$ iff there exists $k' \in [k+a, k+b]$ such that

$$(\mathbf{x}, k') \models \Phi_2 \wedge (\forall k'' \in [k, k'])(\mathbf{x}, k'') \models \Phi_1,$$

i.e., Φ_2 will eventually be satisfied at some instant between $[k+a, k+b]$ and Φ_1 holds consistently before then. Furthermore, we can also induce temporal operators:

- “eventually” $\mathbf{F}_{[a,b]}\Phi := \top \mathbf{U}_{[a,b]}\Phi$ such that it holds when $(\mathbf{x}, k) \models \Phi$ for some $k' \in [k+a, k+b]$;
- “always” $\mathbf{G}_{[a,b]}\Phi := \neg \mathbf{F}_{[a,b]}\neg \Phi$ such that it holds when $(\mathbf{x}, k) \models \Phi$ for any $k' \in [k+a, k+b]$.

We write $\mathbf{x} \models \Phi$ whenever $(\mathbf{x}, 0) \models \Phi$. We assume that the operation horizon T is sufficiently long to evaluate the satisfaction of Φ .

3. Problem formulation

3.1. Fragment of STL formulae

In this paper, we consider the following slightly restricted fragments of STL formulae

$$\varphi ::= \top \mid \pi^\mu \mid \neg \varphi \mid \varphi_1 \wedge \varphi_2 \quad (2a)$$

$$\Phi ::= \mathbf{F}_{[a,b]}\Phi \mid \mathbf{G}_{[a,b]}\Phi \mid \Phi_1 \mathbf{U}_{[a,b]}\Phi_2 \mid \Phi_1 \wedge \Phi_2 \mid \varphi. \quad (2b)$$

Compared to the full fragment of STL, we require that negation can only be applied to Boolean operators. Therefore, the overall STL formula considered is the conjunction of a set of sub-formulae in which nested temporal operators can be applied to an arbitrary Boolean formula.

For technical purposes, we introduce a new temporal operator \mathbf{U}' , defined by: $(\mathbf{x}, k) \models \Phi_1 \mathbf{U}'_{[a,b]}\Phi_2$ iff there exists $k' \in [k+a, k+b]$ such that

$$(\mathbf{x}, k') \models \Phi_2 \wedge (\forall k'' \in [k+a, k'])(\mathbf{x}, k'') \models \Phi_1.$$

Compared to the original definition of “until”, the key difference is that the effective horizon of $\mathbf{U}_{[a,b]}$ is $[0, b]$, while the effective horizon of $\mathbf{U}'_{[a,b]}$ is $[a, b]$. Throughout this paper, we will refer to $\mathbf{U}'_{[a,b]}$ as the “until” operator. This replacement is primarily technical and does not lose generality, as the standard \mathbf{U} can always be expressed as

$$(\mathbf{x}, k) \models \Phi_1 \mathbf{U}_{[a,b]}\Phi_2 \Leftrightarrow (\mathbf{x}, k) \models (\Phi_1 \mathbf{U}'_{[a,b]}\Phi_2) \wedge (\mathbf{G}_{[0,a]}\Phi_1).$$

Recall that for the predicate π^μ in (2a), its satisfaction region, denoted by \mathcal{H}^μ , is the solution to the inequality $\mu(x) \geq 0$, i.e., $\mathcal{H}^\mu = \{x \in \mathcal{X} \mid \mu(x) \geq 0\}$. For other Boolean operators, we have $\mathcal{H}^{\neg\varphi} = \mathcal{X} \setminus \mathcal{H}^\varphi$ and $\mathcal{H}^{\varphi_1 \wedge \varphi_2} = \mathcal{H}^{\varphi_1} \cap \mathcal{H}^{\varphi_2}$. The satisfaction region of constant true predicate \top is the entire state space \mathcal{X} . Therefore, instead of writing φ , we will hereafter simply denote it as $x \in \mathcal{H}^\mu$ or $x \in \mathcal{H}$, using its satisfaction region.

Additionally, while we consider the temporal operator “eventually” (\mathbf{F}) in the semantics, it is subsumed by “until” (\mathbf{U}') since $\mathbf{F}_{[a,b]}\Phi$ can be expressed as $x \in \mathcal{X} \mathbf{U}'_{[a,b]}\Phi$. Thus, we only need to handle the temporal operators \mathbf{G} and \mathbf{U}' from a technical standpoint.

Based on the above discussion, the STL formula Φ in (2) can be equivalently expressed as:

$$\Phi ::= \mathbf{G}_{[a,b]}\Phi \mid \Phi_1 \mathbf{U}'_{[a,b]}\Phi_2 \mid \Phi_1 \wedge \Phi_2 \mid x \in \mathcal{H} \quad (3)$$

3.2. Online monitoring of STL

At time $k \leq T$, the system only generates a partial signal $\mathbf{x}_{0:k} = x_0 x_1 \dots x_k$ (called a prefix) and the remaining signals $\mathbf{x}_{k+1:T}$ (called suffix) will be available in the future. Therefore, with a partial signal $\mathbf{x}_{0:k}$, we denote by

- $\mathbf{x}_{0:k} \models \Phi$ if $\forall \mathbf{x}_{k+1:T} \in \mathcal{X}^{T-k} : \mathbf{x}_{0:k} \mathbf{x}_{k+1:T} \models \Phi$;
- $\mathbf{x}_{0:k} \not\models \Phi$ if $\forall \mathbf{x}_{k+1:T} \in \mathcal{X}^{T-k} : \mathbf{x}_{0:k} \mathbf{x}_{k+1:T} \not\models \Phi$;
- $\mathbf{x}_{0:k} \models_{\gamma} \Phi$ otherwise.

Note that the above partial signal evaluation does not account for system dynamics. In other words, some suffixes in $\mathbf{x}_{k+1:T} \in \mathcal{X}^{T-k}$ may be dynamically infeasible. Thus, in the context of *model-predictive monitoring*, we further classify a prefix signal $\mathbf{x}_{0:k}$ as

- **violated** if, for any control input $\mathbf{u}_{k:T-1}$, we have $\mathbf{x}_{0:k} \xi_f(x_k, \mathbf{u}_{k:T-1}) \not\models \Phi$;
- **feasible** if, for some control input $\mathbf{u}_{k:T-1}$, we have $\mathbf{x}_{0:k} \xi_f(x_k, \mathbf{u}_{k:T-1}) \models \Phi$.

Intuitively, a prefix signal is violated when either the current state already violates the specification or when all possible future trajectories will violate it inevitably. For instance, consider a safety specification $\mathbf{G}_{[0,T]}x \in \mathcal{H}$. If the system reaches any state $x_k \notin \mathcal{H}$ for $k < T$, this immediately constitutes a violation. More subtly, even when the current state satisfies the predicate, the prefix is violated if from state x_k there exists no control sequence $\mathbf{u}_{k:T-1}$ that can generate a trajectory $\xi_f(x_k, \mathbf{u}_{k:T-1})$ remaining entirely within \mathcal{H} throughout the remaining horizon.

Problem 1 (Model-Predictive Online Monitoring). Given a system with dynamic in (1) and an STL formula Φ , design an online monitor

$$\mathcal{M} : \mathcal{X}^* \rightarrow \{\text{vio}, \text{feas}\} \quad (4)$$

such that, for any prefix $\mathbf{x}_{0:k} \in \mathcal{X}^*$, we have

- $\mathcal{M}(\mathbf{x}_{0:k}) = \text{vio}$ iff $\mathbf{x}_{0:k}$ is violated;
- $\mathcal{M}(\mathbf{x}_{0:k}) = \text{feas}$ iff $\mathbf{x}_{0:k}$ is feasible.

4. Tree of nested STL formulae

The main challenge in handling STL formulae with nested temporal operators lies in the dependency between inner and outer operators. For instance, consider the STL formula $\mathbf{G}_{[a_1, b_1]} \mathbf{F}_{[a_2, b_2]} \varphi$. This formula requires that the inner formula $\mathbf{F}_{[a_2, b_2]} \varphi$ must be satisfied for all time instants between a_1 and b_1 . Specifically, starting from any time instant within $[a_1, b_1]$, the system should satisfy φ within the interval $[a_2, b_2]$ from *that point onward*. Consequently, evaluating the correctness of a trajectory requires a total horizon of $[a_1 + a_2, b_1 + b_2]$. To address this dependency clearly, we introduce the concepts of *syntax tree* and *satisfaction vector*. These tools help systematically resolve the dependencies between nested temporal operators.

4.1. Syntax trees for STL formulae

Hereafter, we will equivalently represent an STL formula by a *rooted tree*, which is a finite directed acyclic graph (DAG) with a unique root node and no cycles. It can be expressed as a 4-tuple

$$\mathcal{T} = (V, L, E, v_{\text{root}}),$$

where:

- $V = \{v_1, v_2, \dots, v_m\}$ is the set of nodes;
- $L = \{l, r, \perp\}$ is a label set, where l and r denote “left” and “right”, respectively, and \perp denotes “no order”;
- $E \subseteq V \times L \times V$ is the set of edges, where each $(v, \ell, v') \in E$ indicates that v' is a child node of v , and is specifically a left or right child if $\ell = l$ or $\ell = r$, respectively;
- $v_{\text{root}} \in V$ is the unique root node with no parent.

Clearly, a rooted tree induces a partial order $<$ on V , where $v' < v$ iff v is an ancestor of v' . For each node v , we denote by $\text{child}(v) = \{v' \mid \exists \ell \in L : (v, \ell, v') \in E\}$ the set of its children. Also, we denote by $\text{child-l}(v)$ and $\text{child-r}(v)$ the left and right child of v , respectively, if they exist.

In order to represent an STL formula as a tree, we introduce the following four types of nodes:

- **H-nodes:** Each node represents a satisfaction region of a Boolean formula. Such nodes have no children, as Boolean formulae are always at the innermost level of the entire formula.
- **\wedge -nodes:** Each node represents the Boolean operator “conjunction” and has two or more unordered children. These nodes are evaluated instantly, meaning there is no time interval associated with them.
- **G-nodes:** Each node represents the temporal operator “always” and has exactly one child. Such nodes are associated with a time interval that determines the evaluation period of their descendants.
- **U'-nodes:** Each node represents the temporal operator “until” and has both a left and a right child. These nodes are also associated with a time interval that determines the evaluation period of their descendants.

Now, we are ready to formally define the syntax tree.

Definition 1 (Syntax Trees). Let Φ be an STL formula defined by syntax in (3). The syntax tree of STL formula Φ , denoted by $\mathcal{T}_\Phi = (V_\Phi, L, E_\Phi, v_{\text{root}, \Phi})$, is defined recursively as follows:

- If $\Phi = \Phi_1 \wedge \Phi_2 \wedge \dots \wedge \Phi_n$, then $V_\Phi = (\cup_{i=1}^n V_{\Phi_i}) \cup \{v_\wedge\}$, $E_\Phi = (\cup_{i=1}^n E_{\Phi_i}) \cup \{(v_\wedge, \perp, v_{\text{root}, \Phi_i}) \mid i = 1, \dots, n\}$, and $v_{\text{root}, \Phi} = v_\wedge$, where v_\wedge is a new \wedge -node that does not exist in each \mathcal{T}_{Φ_i} ;
- If $\Phi = \mathbf{G}_{[a, b]} \Phi'$, then $V_\Phi = V_{\Phi'} \cup \{v_G\}$, $E_\Phi = E_{\Phi'} \cup \{(v_G, \perp, v_{\text{root}, \Phi'})\}$, and $v_{\text{root}, \Phi} = v_G$, where v_G is a new G-node that does not exist in $\mathcal{T}_{\Phi'}$ and it is associated with time interval $[a, b]$;
- If $\Phi = \Phi_1 \mathbf{U}'_{[a, b]} \Phi_2$, then $V_\Phi = V_{\Phi_1} \cup V_{\Phi_2} \cup \{v_{U'}\}$, $E_\Phi = E_{\Phi_1} \cup E_{\Phi_2} \cup \{(v_{U'}, l, v_{\text{root}, \Phi_1}), (v_{U'}, r, v_{\text{root}, \Phi_2})\}$, and $v_{\text{root}, \Phi} = v_{U'}$, where $v_{U'}$ is a new U'-node does not exist in \mathcal{T}_{Φ_1} or \mathcal{T}_{Φ_2} , and it is also associated with time interval $[a, b]$;
- If $\Phi = x \in \mathcal{H}$, then \mathcal{T}_Φ only contains a single H-node associated with predicate $x \in \mathcal{H}$.

For each node $v \in V_\Phi$, we denote by $[a^v, b^v]$ its **associated time interval**, with $a^v = b^v = 0$ when v is a \wedge -node or an H-node. Also, for each H-node $v \in V_{\mathcal{H}}$, we denote by H^v the associated predicate region. Hereafter, we will omit the subscript Φ and refer to $\mathcal{T} = (V, L, E, v_{\text{root}})$ as the syntax tree of the formula Φ when the context is clear. We use \mathcal{T}^v to denote the subtree with v as the root node. The subtree \mathcal{T}^v represents a sub-formula Φ^v of the original STL formula Φ . For $\star \in \{\wedge, \mathbf{G}, \mathbf{U}', \mathcal{H}\}$, we denote by $V_\star \subseteq V$ the set of all \star -nodes in \mathcal{T} . Furthermore, let

$$V = \{v_1, v_2, \dots, v_m\},$$

and we assume that all H-nodes are ordered as the first $h \geq 1$ elements, i.e., $V_{\mathcal{H}} = \{v_1, \dots, v_h\}$.




Fig. 1. Syntax Tree of STL formula (5).

4.2. Satisfaction vector

As discussed, for each node $v \in V$ in the syntax tree, it evaluates its descendant nodes within the time interval $[a^v, b^v]$. Note that this interval is relative to the local perspective, as v may have ancestor nodes in the tree corresponding to outer operators whose time intervals also apply to v . Therefore, to determine the absolute time interval for evaluating node v from a global perspective, one must consider the local time intervals of all its ancestors.

Formally, for each node $v \in V$, we denote by $\text{ances}(v) = \{v' \in V \mid v < v'\}$ the set of all its ancestors. This set includes all nodes on the unique path from the root node to v in \mathcal{T} , excluding v itself. Based on this, we introduce the following definition of *evaluation horizon*.

Definition 2 (Evaluation Horizons). Let \mathcal{T} be the syntax tree of STL formula Φ and $v \in V$ be a node. Then the *evaluation horizon* of node v is defined by

$$[\text{int}^v, \text{end}^v] = [\sum_{v' \in \text{ances}(v)} a^{v'}, \sum_{v' \in \text{ances}(v)} b^{v'}].$$

For example, consider the following STL formula

$$\Phi = ((G_{[1,3]}x \in H_1)U'_{[2,5]}x \in H_2) \wedge (G_{[3,7]}x \in H_3). \tag{5}$$

Its syntax tree \mathcal{T} is shown in Fig. 1. For H node v_1 , its evaluation horizon is $[\text{int}^{v_1}, \text{end}^{v_1}] = [2 + 1 = 3, 5 + 3 = 8]$.

Therefore, for each time instant $t \in [\text{int}^v, \text{end}^v]$, it is meaningful to discuss the satisfaction status of node v . In general, there are three possible satisfaction statuses for a node: (i) **satisfied**, denoted by “1”; (ii) **violated**, denoted by “0”; and (iii) **uncertain**, denoted by “?”. The last case arises when there is insufficient information to evaluate the node. For example, consider the formula $G_{[10,12]}x \in H$. If the current time instant is $k = 8$, then the satisfaction status of the H -node for $t = 10, 11, 12$ is uncertain because the evaluation horizon of the H -node has not yet been reached. However, if at time instant $k = 10$, the current state x_k falls within the region H , then the satisfaction status of the H -node at $t = 10$ becomes 1. However, the satisfaction status for future instants $t = 11, 12$ remains uncertain.

To capture this, we define the satisfaction status for each time instant within the evaluation horizon of a node as the *satisfaction vector*, as follows.

Definition 3 (Satisfaction Vectors). Let \mathcal{T} be the syntax tree of STL formula Φ and $v \in V$ be a node with evaluation horizon $[\text{int}^v, \text{end}^v]$. Then a *satisfaction vector* of node v is a vector of form

$$t = (t[\text{int}^v], t[\text{int}^v + 1], \dots, t[\text{end}^v]) \in \{0, 1, ?\}^H,$$

where $H = \text{end}^v - \text{int}^v + 1$ is the length of its evaluation horizon and $t[t], t = \text{int}^v, \dots, \text{end}^v$ denotes its t th element.

For example, consider the leaf node v_1 associated with $x \in H_1$ in Formula (5). Its evaluation horizon is $[3, 8]$. A specific satisfaction vector for v_1 could be $t^{v_1} = (1, 1, 0, ?, ?, ?)$ when $t = 6$, where the length is $8 - 3 + 1 = 6$. This vector indicates that:

- At time $t = 3$ and $t = 4$, the system state is inside H_1 (status 1);
- At time $t = 5$, the state is outside H_1 (status 0);
- For $t = 6, 7, 8$, the status is currently unknown (?).

This concrete vector forms part of the basic satisfaction vector tuple used in our algorithm.

In the above definition, we count the first element of t starting from $t[\text{int}^v]$ rather than $t[1]$ to align with the absolute time instant from the perspective of the root node. This ensures consistency between the indexing of the satisfaction vector and the actual time instants being evaluated.

For each node, its satisfaction vector is essentially determined by the vectors of its children nodes. Through an inductive argument based on the tree structure, once the satisfaction vector of each \mathcal{H} -node is known, we can compute the satisfaction vectors of all non-leaf nodes. Therefore, we refer to the satisfaction vector of an \mathcal{H} -node as a *basic vector*, as it serves as the foundation for *inducing vectors* for other nodes. This leads to the following definitions.

Definition 4 (Basic Vectors). A *basic set* of satisfaction vectors is a tuple of satisfaction vectors of the form

$$I = (i^{v_1}, \dots, i^{v_h}) \in \{0, 1, ?\}^{H_1} \times \dots \times \{0, 1, ?\}^{H_h},$$

where $V_{\mathcal{H}} = \{v_1, \dots, v_h\}$ is the set of all \mathcal{H} -nodes and for each $i = 1, \dots, h$, $H_i = \text{end}^{v_i} - \text{int}^{v_i} + 1$ represents the length of the evaluation horizon of v_i .

Hereafter, we will maintain basic vectors as time-updated information. For each time instant $k = 0, 1, \dots, T$, a basic set $I_k = (i_k^{v_1}, \dots, i_k^{v_h})$ must satisfy the following condition for each $v \in V_{\mathcal{H}}$ and $t = \text{int}^v, \dots, \text{end}^v$:

- If $t \geq k$, then $i_k^v[t] = ?$ (unknown status for futures);
- If $t < k$, then $i_k^v[t] \neq ?$ (determined status for pasts).

This reflects the constraint that satisfaction status cannot be evaluated for future time instants, while for past instants the status must be resolved to either 0 (violated) or 1 (satisfied). We denote by \mathcal{I}_k the set of all potential basic vectors at instant k satisfying the above two constraints.

Definition 5 (Induced Vectors). Given a basic set of satisfaction vectors $I = (i^{v_1}, \dots, i^{v_h})$, the *induced set of satisfaction vectors* for the remaining nodes, denoted by $I^+ = (i^{v_{h+1}}, \dots, i^{v_m})$, is defined recursively by: for each $v \in \{v_{h+1}, \dots, v_m\}$, we have

- If $v \in V_{\wedge}$, then for each $t = \text{int}^v, \dots, \text{end}^v$, we have

$$i^v[t] = \begin{cases} 0 & \text{if } \exists v' \in \text{child}(v) : i^{v'}[t] = 0 \\ 1 & \text{if } \forall v' \in \text{child}(v) : i^{v'}[t] = 1 \\ ? & \text{otherwise} \end{cases} .$$

- If $v \in V_G$, then for each $t = \text{int}^v, \dots, \text{end}^v$, we have

$$i^v[t] = \begin{cases} 0 & \text{if } \exists t' \in [a^v, b^v] : i^{v'}[t+t'] = 0 \\ 1 & \text{if } \forall t' \in [a^v, b^v] : i^{v'}[t+t'] = 1 \\ ? & \text{otherwise} \end{cases} .$$

where v' is the unique child node of v .

- If $v \in V_{\vee}$, then for each $t = \text{int}^v, \dots, \text{end}^v$, we have

$$i^v[t] = \begin{cases} 0 & \text{if } \left[\begin{array}{l} (\forall t' \in [a^v, b^v])(i^{v_l}[t+t'] = 0) \\ \text{or} \\ (\forall t' \in [a^v, b^v] : i^{v_r}[t+t'] = 1) \\ (\exists t'' \in [a^v, t']) (i^{v_l}[t+t''] = 0) \end{array} \right] , \\ 1 & \text{if } \left[\begin{array}{l} (\exists t' \in [a^v, b^v] : i^{v_r}[t+t'] = 1) \\ (\forall t'' \in [a^v, t']) (i^{v_l}[t+t''] = 1) \end{array} \right] \\ ? & \text{otherwise} \end{cases} ,$$

where v_l and v_r denotes the left and right child nodes of v , respectively.

The intuition behind the above definition is essentially the semantics of STL formulae. The induced vectors are well-defined because all nodes are ordered in the tree structure, where the basic vectors representing \mathcal{H} -nodes are leaves in the tree with no children. Hence, given a basic set of satisfaction vectors, one can compute its induced vectors in a bottom-up manner, starting from the leaf nodes and progressing to the root node.

For example, let v_G be the parent node of v_1 , representing the sub-formula $\mathbf{G}_{[1,3]}x \in \mathcal{H}_1$. Its evaluation horizon is $[2, 5]$. The induced vector i^{v_G} is derived from the basic vector i^{v_1} of its child. For instance, to determine the status of i^{v_G} at time $t = 2$, we check if $i^{v_1}[2 + \tau] = 1$ for all $\tau \in [1, 3]$. If i^{v_1} has 1 at indices 3, 4, 5, then $i^{v_G}[2]$ resolves to 1. Otherwise, if any entry is 0, it resolves to 0. This recursive “bottom-up” calculation allows us to determine the status of the root node from the basic vectors.

5. Online monitoring algorithms

5.1. Evolution of basic vectors

To track the progress of the STL formulae without storing the entire state trajectory, our approach is to maintain only the basic satisfaction vectors and update them recursively upon observing the new system state at each time instant. This recursive computation proceeds as follows:

- **Initialization:** At time instant $k = 0$, prior to observing any system state, the initial basic set is defined as:

$$I_0 = (i_0^{v_1}, i_0^{v_2}, \dots, i_0^{v_h}),$$

where each $i_0^{v_i} = (?, \dots, ?)$ is a vector with all entries initialized to the uncertain status “?”. This reflects complete uncertainty about future satisfaction before any state observations are made.

- **Online Update:** At time instant k , given the current basic satisfaction vector set $I_k = (i_k^{v_1}, \dots, i_k^{v_h})$, the update upon observing a new state x_k is defined as:

$$I_{k+1} = (i_{k+1}^{v_1}, \dots, i_{k+1}^{v_h}) = \text{update}(I_k, x_k),$$

where for each $v_i \in \{v_1, \dots, v_h\}$ and $t \in [\text{int}^{v_i}, \text{end}^{v_i}]$, the updated satisfaction vector entry is:

$$i_{k+1}^{v_i}[t] = \begin{cases} i_k^{v_i}[t] & \text{if } t \neq k, \\ 0 & \text{if } t = k \wedge x_k \notin \mathcal{H}^{v_i}, \\ 1 & \text{if } t = k \wedge x_k \in \mathcal{H}^{v_i}. \end{cases} \quad (6)$$

Therefore, the uncertain status “?” in the basic vectors is resolved at time k based on state x_k .

Let $\mathbf{x}_{0:k-1} = x_0 x_1 \dots x_{k-1}$ be a state sequence of length k . We denote by $I(\mathbf{x}_{0:k-1})$ the basic set of satisfaction vectors reached recursively by the state sequence $\mathbf{x}_{0:k-1}$ from I_0 . On the other hand, for any basic set I , we define the set of state sequences of length k consistent with I as

$$\mathbf{x}_{0:k-1}^I = \{\mathbf{x}_{0:k-1} \in \mathcal{X}^k \mid I(\mathbf{x}_{0:k-1}) = I\}.$$

The following result demonstrates that the recursive computation of the basic set above indeed captures all task-relevant information from the complete state trajectory.

Proposition 1. For any state sequence $\mathbf{x}_{0:k}$ and sub-formula Φ^v we have

$$(\mathbf{x}_{0:k}, t) \models \Phi^v \Leftrightarrow i_{I(\mathbf{x}_{0:k})}^{v_i}[t] = 1.$$

Proof. Prove this proposition recursively by mathematical induction. For the case of leaf node $\Phi^v = x \in \mathcal{H}$:

\Rightarrow : When $(\mathbf{x}_{0:k}, t) \models \Phi^v$, there must be $x_t \in \mathcal{H}$. According to (6), $i_{I(\mathbf{x}_{0:k})}^{v_i}[t] = 1$.

\Leftarrow : When $i_{I(\mathbf{x}_{0:k})}^{v_i}[t] = 1$, which also means $x_t \in \mathcal{H}$. So $(\mathbf{x}_{0:k}, t) \models \Phi^v$.

The proposition holds for leaf nodes. Assuming it always holds for child nodes, now prove the case of other nodes.

(1) $\Phi^v = \Phi^{v_1^c} \wedge \Phi^{v_2^c} \wedge \dots \wedge \Phi^{v_n^c}$:

\Rightarrow : When $(\mathbf{x}_{0:k}, t) \models \Phi^v$, there must be $(\mathbf{x}_{0:k}, t) \models \Phi^{v_1^c}, (\mathbf{x}_{0:k}, t) \models \Phi^{v_2^c}, \dots, (\mathbf{x}_{0:k}, t) \models \Phi^{v_n^c}$. So $\forall v_i^c \in \text{child}(v)$, $i_{I(\mathbf{x}_{0:k})}^{v_i^c}[t] = 1$. Then $i_{I(\mathbf{x}_{0:k})}^{v_i}[t] = 1$ according to Definition 5.

\Leftarrow : Reverse the proof of \Rightarrow .

(2) $\Phi^v = \mathbf{G}_{[a,b]} \Phi^{v^c}$:

\Rightarrow : When $(\mathbf{x}_{0:k}, t) \models \Phi^v$, there must be $\forall t' \in [t+a, t+b], (\mathbf{x}_{0:k}, t') \models \Phi^{v^c}$. So $\forall t' \in [t+a, t+b], i_{I(\mathbf{x}_{0:k})}^{v_i^c}[t'] = 1$. So $\forall v_i^c \in \text{child}(v)$, $i_{I(\mathbf{x}_{0:k})}^{v_i^c}[t] = 1$. Then $i_{I(\mathbf{x}_{0:k})}^{v_i}[t] = 1$ according to Definition 5.

\Leftarrow : Reverse the proof of \Rightarrow .

(3) $\Phi^v = \Phi^{v^c} \mathbf{U}'_{[a,b]} \Phi^{v'^c}$:

\Rightarrow : When $(\mathbf{x}_{0:k}, t) \models \Phi^v$, there must be $\exists t' \in [t+a, t+b], (\mathbf{x}_{0:k}, t') \models \Phi^{v'^c}$ i.e. $i_{I(\mathbf{x}_{0:k})}^{v_i'^c}[t'] = 1$ and $\forall t'' \in [t+a, t']$ such that $(\mathbf{x}_{0:k}, t'') \models \Phi^{v^c}$ i.e. $i_{I(\mathbf{x}_{0:k})}^{v_i^c}[t''] = 1$. So $i_{I(\mathbf{x}_{0:k})}^{v_i}[t] = 1$ according to Definition 5.

\Leftarrow : Reverse the proof of \Rightarrow .

According to the above, we prove $(\mathbf{x}_{0:k}, t) \models \Phi^v \Leftrightarrow i_{I(\mathbf{x}_{0:k})}^{v_i}[t] = 1$. \square

Given a basic set of satisfaction vectors I , its induced set of satisfaction vectors is defined as $I^+ = (i^{v_{h+1}}, \dots, i^{v_m})$. We call the satisfaction vector corresponding to the root node v_{root} in I^+ the *root satisfaction vector*, denoted by $i_{I^+}^{\text{root}} \in I^+$. The following result demonstrates that this root vector fully captures the satisfaction status of the entire STL formula.

Proposition 2. For any state sequence $\mathbf{x}_{0:k}$, we have

$$\mathbf{x}_{0:k} \models \Phi \Leftrightarrow i_{I(\mathbf{x}_{0:k})}^{\text{root}}[0] = 1. \quad (7)$$

Proof. As a special case of [Proposition 1](#), when $v = v_{\text{root}}$ and $t = 0$, $\mathbf{x}_{0:k} \models \Phi \Leftrightarrow I_{I(\mathbf{x}_{0:k})}^{\text{root}}[0] = 1$, which means the root vector fully captures the satisfaction status of the entire STL formula. \square

For simplicity, we will write $I_{I(\mathbf{x}_{0:k})}^{\text{root}} = 0, 1, ?$ whenever $I_{I(\mathbf{x}_{0:k})}^{\text{root}}[0] = 0, 1, ?$ as the root vector only has one element.

Proposition 3. Let I be a basic set and k be a time instant. For any two sequences $\mathbf{x}'_{0:k-1}, \mathbf{x}''_{0:k-1} \in \mathbf{x}^I_{0:k-1}$ consistent with I , and any future sequence $\mathbf{x}_{k:T} = x_k x_{k+1} \dots x_T$, we have

$$\mathbf{x}'_{0:k-1} \mathbf{x}_{k:T} \models \Phi \Leftrightarrow \mathbf{x}''_{0:k-1} \mathbf{x}_{k:T} \models \Phi. \tag{8}$$

Proof. According to [Proposition 2](#), we know that

$\mathbf{x}'_{0:k-1} \mathbf{x}_{k:T} \models \Phi \Leftrightarrow I_{I(\mathbf{x}'_{0:k-1} \mathbf{x}_{k:T})}^{\text{root}}[0] = 1$. For $\mathbf{x}'_{0:k-1}, \mathbf{x}''_{0:k-1} \in \mathbf{x}^I_{0:k-1}$, we have $I(\mathbf{x}'_{0:k-1} \mathbf{x}_{k:T}) = I(\mathbf{x}''_{0:k-1} \mathbf{x}_{k:T})$. So there must be $I_{I(\mathbf{x}'_{0:k-1} \mathbf{x}_{k:T})}^{\text{root}}[0] = I_{I(\mathbf{x}''_{0:k-1} \mathbf{x}_{k:T})}^{\text{root}}[0] = 1$. Then due to [Proposition 2](#), we have $\mathbf{x}''_{0:k-1} \mathbf{x}_{k:T} \models \Phi$. The proof above also holds in reverse. So there is $\mathbf{x}'_{0:k-1} \mathbf{x}_{k:T} \models \Phi \Leftrightarrow \mathbf{x}''_{0:k-1} \mathbf{x}_{k:T} \models \Phi$. \square

Based on the above result, it is meaningful to write

$$\mathbf{x}^I_{0:k-1} \mathbf{x}_{k:T} \models \Phi \tag{9}$$

whenever $\mathbf{x}'_{0:k-1} \mathbf{x}_{k:T} \models \Phi$ holds for some (or equivalently, for all, due to [Proposition 3](#)) $\mathbf{x}'_{0:k-1} \in \mathbf{x}^I_{0:k-1}$.

5.2. Model-predictive monitor

By recursively maintaining the basic set $I_k = I(\mathbf{x}_{0:k-1})$, we can draw the following conclusions about the current status of the entire specification formula:

- If $I_k^{\text{root}} = 0$, the specification has been violated, and the monitor should output $\mathcal{M}(\mathbf{x}_{0:k-1}) = \text{vio}$;
- If $I_k^{\text{root}} = 1$, the specification has been satisfied, and no further monitoring is required.

These conclusions are independent of the system dynamics, as the satisfaction status is fully determined in these cases. However, when $I_k^{\text{root}} = ?$, both vio and feas remain possible outcomes, requiring analysis of the system dynamics. This leads to the following key definition.

Definition 6 (I-Determined Feasible Sets). Let Φ be an STL formula, $k \in [0, T]$ be a time instant and I be a basic set. Then the *I-determined feasible set* at instant k , denoted by $X^I_k \subseteq \mathcal{X}$, is the set of states from which there exists a solution $\mathbf{u}_{k:T-1}$ that satisfies Φ given the current basic set I , i.e.,

$$X^I_k = \left\{ x_k \in \mathcal{X} \mid \begin{array}{l} \exists \mathbf{u}_{k:T-1} \in \mathcal{U}^{T-k} \\ \text{s.t. } \mathbf{x}'_{0:k-1} x_k \xi_f(x_k, \mathbf{u}_{k:T-1}) \models \Phi \end{array} \right\} \tag{10}$$

Based on the above notion, now we present our main online monitoring algorithm as shown in [Algorithm 1](#). The algorithm initializes with the time instant $k = 0$ and the initial basic set $I = I_0$ (line 1). The monitoring process iterates as long as the STL formula remains unresolved (line 2). At each iteration, the current state x_k is read (line 3) and checked against the feasible set X^I_k . If $x_k \notin X^I_k$, the algorithm immediately terminates with a violation decision (lines 4–6). Otherwise, the monitoring decision is set to $\mathcal{M} = \text{feas}$, and the basic set I is recursively updated by $\text{update}(I, x_k)$ (lines 7–9). The time instant is then increased (line 10), and the loop continues until either a violation is detected or the STL formula is satisfied.

The correctness of [Algorithm 1](#) is established as follows.

Theorem 1. The online monitor \mathcal{M} defined by [Algorithm 1](#) indeed solves [Problem 1](#).

5.3. Offline computation of feasible sets

The proposed online monitoring algorithm utilizes pre-computed all feasible sets X^I_k at each time instant. This subsection details the offline computation of these sets.

Note that, at each time instant k , we only need to consider those basic vectors for which the entire STL formula Φ has not been violated. Therefore, we define

$$\mathbb{I}_k = \{ I \in \mathcal{I}_k \mid I^{\text{root}} \neq 0 \}$$

the set of **feasible basic vectors** at time instant k .

Furthermore, the basic vector set cannot be updated arbitrarily in the next time instant as it should be consistent with the existing history. Therefore, let $I = (i^{v_1}, \dots, i^{v_h}) \in \mathbb{I}_k$ be a basic vector set for time instant k and $I_s = (i_s^{v_1}, \dots, i_s^{v_h}) \in \mathbb{I}_{k+1}$ be a basic vector set for time instant $k + 1$. We say $I_s \in \mathbb{I}_{k+1}$ is a **successor basic set** of $I \in \mathbb{I}_k$ at instant k if

$$\forall v \in V_{\mathcal{H}}, \forall t \in [\text{int}^v, \min(k-1, \text{end}^v)] : i^v[t] = i_s^v[t].$$

Algorithm 1: Online Monitoring algorithm

Input: Feasible sets X_k^I computed offline
Output: Online monitoring decision $\mathcal{M}(x_{0:k})$

```

1  $k \leftarrow 0, I \leftarrow I_0$ 
2 while  $t_I^{\text{root}} \neq 1$  do
3   observe a new current state  $x_k$ 
4   if  $x_k \notin X_k^I$  then
5      $\mathcal{M}$  issues “vio”
6     return “ $\Phi$  is violated”
7   else
8      $\mathcal{M}$  issues “feas”
9      $I \leftarrow \text{update}(I, x_k)$ 
10   $k \leftarrow k + 1$ 

```

We denote by $\text{succ}(I, k) \subseteq \mathbb{I}_{k+1}$ the set of all successor basic sets of I at instant k . Moreover, when the basic set evolves from I to I_s at instant k , it must reach a state x_k consistent to the update rule. We define the region that x_k needs to be in as *consistent region*.

Definition 7 (Consistent Regions). Let $I = (t^{v_1}, \dots, t^{v_h}) \in \mathbb{I}_k$ be a basic vector at instant k , and $I_s = (t_s^{v_1}, \dots, t_s^{v_h}) \in \text{succ}(I, k)$ be a successor basic set of I . In order to trigger the evolution of the basic set from I to I_s at instant k , the system should be in the consistent region $H_k(I, I_s)$, which is defined by

$$H_k(I, I_s) = \bigcap_{i=1}^h H_{ik}, \quad (11)$$

where

$$H_{ik} = \begin{cases} \mathcal{X} \setminus \mathcal{H}^{v_i} & \text{if } t_s^{v_i}[k] = 0 \\ \mathcal{H}^{v_i} & \text{if } t_s^{v_i}[k] = 1 \\ \mathcal{X} & \text{if } k > \text{end}^{v_i} \text{ or } k < \text{int}^{v_i} \end{cases}$$

Finally, in order to ensure that the prefix signal is always feasible, the feasible set X_k^I needs to be able to reach region $X_{k+1}^{I_s}$ in one step. This is formalized by the *one-step feasible set* defined as follows.

Definition 8 (One-Step Feasible Set). Let $S \subseteq \mathcal{X}$ be a set of states representing the “target region”. Then the one-step feasible set of S is defined by

$$Y(S) = \{x \in \mathcal{X} \mid \exists u \in \mathcal{U} \text{ s.t. } f(x, u) \in S\}. \quad (12)$$

We define $\mathbb{X}_k = \{X_k^I \mid I \in \mathbb{I}_k\}$ as the set of all feasible sets for instant k . Then our offline objective is to compute all possible $\mathbb{X}_0, \mathbb{X}_1, \dots, \mathbb{X}_T$, which are used as a look-up table during the online monitoring process. In terms of our computation of feasible regions, if the system is evolving from I to I_s and maintains the satisfiability of I_s from instant $k+1$, then we know that the system should be in region $H_k(I, I_s) \cap Y(X_{k+1}^{I_s})$ at instant k . However, the basic set I_s for the next instant depends on the current state of the system. Therefore, to compute X_k^I , we need to consider all possible successor sets $I_s \in \text{succ}(I, k)$, and take the union of these regions. This is formalized by the following equation

$$X_k^I = \bigcup_{I_s \in \text{succ}(I, k)} \left(H_k(I, I_s) \cap Y(X_{k+1}^{I_s}) \right). \quad (13)$$

Based on the above equation, we compute all feasible sets across the entire horizon using a backward recursion procedure, as formalized in Algorithm 2. The algorithm initializes by setting the terminal feasible set $X_{T+1}^I \leftarrow \mathbb{R}^n$ for all $I \in \mathbb{I}_{T+1}$ (lines 3–4). The recursion proceeds backward in time from $k = T$ to $k = 0$. If the STL formula is already satisfied (i.e., $t_I^{\text{root}} = 1$), then set $X_k^I \leftarrow \mathbb{R}^n$. At each time step k , the feasible set X_k^I is computed for every basic set $I \in \mathbb{I}_k$ through two operations:

- Intersection with $H_k(I, I_s)$ and $Y(X_{k+1}^{I_s})$, and
- Union over all successor sets $I_s \in \text{succ}(I, k)$ to ensure completeness.

The resulting sets are aggregated into \mathbb{X}_k , which captures all feasible states at time k (lines 5–13).

Computation of One-Step Feasible Sets: Calculating the exact one-step feasible set $Y(S)$ depends heavily on the system dynamics f . For linear systems with polytopic constraints, this can be computed efficiently using polyhedral operations (e.g., Minkowski sum and projection). In this work, to handle general nonlinear dynamics (as considered in our case studies), we adopt a grid-based abstraction approach. The state space \mathcal{X} and input space \mathcal{U} are discretized into finite sets of grid points, denoted by $\mathcal{X}_{\text{grid}}$ and $\mathcal{U}_{\text{grid}}$,

Algorithm 2: Offline Computations of Feasible Sets

Input: STL formula Φ
Output: All potential sets $\{\mathbb{X}_k : k \in [0, T]\}$

```

1 for  $k \in [0, T]$  do
2    $\mathbb{X}_k \leftarrow \emptyset$ 
3 for  $I \in \mathbb{I}_{T+1}$  do
4    $X_{T+1}^I \leftarrow \mathbb{R}^n$ 
5  $k \leftarrow T$ 
6 while  $k \geq 0$  do
7   for  $I \in \mathbb{I}_k$  do
8     if  $I^{\text{root}} = 1$  then
9        $X_k^I \leftarrow \mathbb{R}^n$ 
10    else
11       $X_k^I \leftarrow \bigcup_{I_s \in \text{succ}(I,k)} (H_k(I, I_s) \cap Y(X_{k+1}^{I_s}))$ 
12     $\mathbb{X}_k \leftarrow \mathbb{X}_k \cup \{X_k^I\}$ 
13   $k \leftarrow k - 1$ 

```

respectively. A state $x \in \mathcal{X}_{grid}$ is included in the approximated one-step feasible set $\hat{Y}(S)$ if there exists a control input $u \in \mathcal{U}_{grid}$ such that the next state $f(x, u)$ falls into a grid cell contained in S . This grid-based approximation allows us to compute the feasible sets for general dynamical systems, albeit with a trade-off between grid resolution and computational complexity.

Complexity Analysis: The computational complexity of the proposed framework is dominated by the offline pre-computation phase, specifically the calculation of one-step feasible sets $Y(\cdot)$ in Algorithm 2. Let n be the number of atomic predicates (leaf nodes) and T be the time horizon. In the worst case, where all possible combinations of satisfaction statuses are feasible, the number of basic vectors to be processed at each time step k is 2^{nk} (considering only 0 and 1 for the current status of each predicate). The total number of $Y(\cdot)$ computations across the entire horizon corresponds to the summation of a geometric series: $\sum_{k=1}^T 2^{nk}$. Consequently, the worst-case offline complexity scales as $O(|\mathcal{X}_{grid}| \cdot 2^{nT})$. While this exponential growth is a known trade-off for handling nested STL specifications, the online monitoring phase remains highly efficient with a complexity of $O(1)$ for set membership checks and vector updates, ensuring suitability for real-time applications.

Comparison with Non-nested Approaches: For non-nested formulae (e.g., $\bigwedge_{i=1}^N \Phi_i$), the approach in [21] maintains an *index set* of remaining unsatisfied sub-formulae. Consequently, its offline augmented state space grows exponentially with the number of parallel sub-formulae N (i.e., 2^N combinations), but does not depend on the length of the time horizon. In contrast, handling nested temporal operators (e.g., $\mathbf{G}_{[0, T_1]} \mathbf{F}_{[0, T_2]} \phi$) requires resolving complex temporal dependencies where the inner formula must be evaluated across multiple future time steps. To achieve this, our method employs *satisfaction vectors* to track the time-wise status of atomic predicates, leading to an offline complexity of $O(|\mathcal{X}_{grid}| \cdot 2^{nT})$. Therefore, compared to [21], the introduction of nested operators fundamentally shifts the exponential complexity dependence from the number of sub-formulae to the length of the evaluation horizons. This increased offline computational overhead is the necessary cost for supporting the significantly richer expressiveness of nested STL specifications, while the online monitoring efficiency remains comparably lightweight.

6. Case studies for online monitoring

In this section, we illustrate our online monitoring algorithm with two cases. We implemented the above methods in Python language.¹

6.1. Building temperatures

We consider the problem of monitoring the temperature of a single zone building whose dynamic is from [27]

$$x_{k+1} = x_k + \tau_s (\alpha_e (T_e - x_k) + \alpha_H (T_h - x_k) u_k),$$

where $x_k \in \mathcal{X} = [0, 45]$ is the zone temperature ($^{\circ}\text{C}$) at step k , $u_k \in \mathcal{U} = [0, 1]$ is the heater valve's normalized position, and $\tau_s = 1$ min is the sampling interval. The model parameters include the heater temperature $T_h = 55^{\circ}\text{C}$, ambient temperature $T_e = 0^{\circ}\text{C}$, and heat transfer coefficients $\alpha_e = 0.06$ (environmental) and $\alpha_H = 0.08$ (heater).

¹ Our codes are available at <https://github.com/sjtu-hantao/MPM4STLnested>, where more details can be found.




Fig. 2. Two trajectories of temperature control system.




Fig. 3. Two trajectories that violate at different instants.

The temperature control system aims to maintain the zone temperature within the comfortable range of 20 °C–25 °C at all times, ensuring that any temperature deviation is corrected within 5 min during any 10-minute window. This requirement can be specified using the STL formula:

$$\Phi = \mathbf{G}_{[0,10]}(\mathbf{F}_{[0,5]}x_k \in [20, 25]).$$

Fig. 2 presents two temperature trajectories that are identical until time step $k = 13$. For the black signal, the condition $x_{13} \in X_{13}^I$ holds, indicating that the control task remains feasible. However, for the red signal, $x_{13} \notin X_{13}^I$, i.e., the specification will inevitably be violated. Hence no control input exists that could bring the temperature within the required range within the time window.

6.2. Autonomous robots

We consider a simple autonomous robot whose dynamic model is given as follows

$$x_{k+1} = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} x_k + \begin{bmatrix} \tau_s & 0 \\ 0 & \tau_s \end{bmatrix} u_k,$$

where state $x_k \in \mathcal{X} = [0, 12] \times [0, 12]$ denotes the position of the robot at instant k , control input $u_k \in \mathcal{U} = [-1, 1] \times [-1, 1]$ is the speed of the robot, and $\tau_s = 1s$ is the sampling time. The objective of the robot is to patrol both region A_1 and region A_2 in 6 s and stay in A_2 for at least 2 s. This task can be described by the following STL formula

$$\Phi = \mathbf{F}_{[0,6]}A_1 \wedge \mathbf{F}_{[0,6]}(\mathbf{G}_{[0,2]}A_2),$$

where $A_1 = (x \in [3, 5]) \wedge (y \in [3, 5])$ and $A_2 = (x \in [6, 8]) \wedge (y \in [6, 8])$.

Let us consider two trajectories of the robot shown in Fig. 3. For the sake of clarity, in each figure, we only draw one feasible set at a certain instant. In the left figure, since $x_0 \notin X_0^I$, the monitor can claim initially that the task cannot be satisfied as it is too far away from the target regions. In the right figure, for $k = 0, 1, 2, 3$, the trajectory stays within the feasible region and the monitor issues “feas”. Yet, we have $x_4 \notin X_4^I$. Therefore, the monitor issues “vio” at $k = 4$.




Fig. 4. Feasible sets of the double integrator model.

6.3. Double integrator with coupled dynamics

We consider a double integrator model describing the longitudinal motion of a vehicle. The state is defined as $x_k = [p_k, v_k]^T$, representing position and velocity, respectively. The discretized dynamics with sampling time $\Delta t = 0.2s$ are given by:

$$x_{k+1} = \underbrace{\begin{bmatrix} 1 & 0.2 \\ 0 & 1 \end{bmatrix}}_A x_k + \underbrace{\begin{bmatrix} 0.02 \\ 0.2 \end{bmatrix}}_B u_k, \tag{14}$$

where the control input $u_k \in [-1, 1]$ represents acceleration. The non-diagonal element $A_{12} = 0.2$ introduces a direct coupling between velocity and position updates. The state space is constrained to $\mathcal{X} = [0, 16] \times [-3, 3]$.

The monitoring task is to ensure that the vehicle reaches a specific target region (position $[8, 10]$ and velocity $[-0.2, 0.2]$) within 17 steps and stays there for at least 3 steps. This nested specification is expressed as:

$$\Phi_{di} = \mathbf{F}_{[0,17]} \mathbf{G}_{[0,3]}(p_k \in [8, 10] \wedge v_k \in [-0.2, 0.2]). \tag{15}$$

The offline computed feasible sets are visualized in Fig. 4. The results clearly show that the feasible sets are not axis-aligned rectangles but sheared shapes. This deformation precisely captures the physical coupling: attaining a specific position requires a velocity constraint that depends on the distance to the target (e.g., higher velocities are only permissible when far from the target). This confirms that our grid-based approach effectively handles coupled dynamics in the calculation of one-step feasible sets.

7. Conclusions

In this paper, we proposed a model-based online monitoring algorithm for nested STL using syntax trees. This structure enables efficient online execution via simple vector updates, though the offline feasible set computation remains demanding for complex dynamics. In future work, we plan to extend this framework in three directions: (1) addressing stochastic systems by integrating probabilistic reachability to output satisfaction probabilities; (2) investigating quantitative monitoring via robust STL semantics to provide richer control information, despite the higher complexity; and (3) exploring syntax tree abstraction techniques to mitigate scalability issues, aiming to reduce the state dimensionality while preserving logical implications.

CRedit authorship contribution statement

Tao Han: Formal analysis, Conceptualization. **Shaoyuan Li:** Supervision, Formal analysis. **Xiang Yin:** Supervision, Formal analysis, Conceptualization.

Declaration of competing interest

The authors declare the following financial interests/personal relationships which may be considered as potential competing interests: Xiang Yin reports financial support was provided by National Natural Science Foundation of China. Xiang Yin reports a relationship with National Natural Science Foundation of China that includes: funding grants. If there are other authors, they declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data availability

No data was used for the research described in the article.

References

- [1] M. Ma, E. Bartocci, E. Lifland, J.A. Stankovic, L. Feng, A novel spatial-temporal specification-based monitoring system for smart cities, *IEEE Internet Things J.* 8 (15) (2021) 11793–11806.
- [2] Y.E. Sahin, R. Quirynen, S. Di Cairano, Autonomous vehicle decision-making and monitoring based on signal temporal logic and mixed-integer programming, in: *American Control Conference*, 2020, pp. 454–459.
- [3] G. Chen, M. Liu, Z. Kong, Temporal-logic-based semantic fault diagnosis with time-series data from industrial internet of things, *IEEE Trans. Ind. Electron.* 68 (5) (2020) 4393–4403.
- [4] E. Bartocci, J. Deshmukh, A. Donzé, G. Fainekos, O. Maler, D. Ničković, S. Sankaranarayanan, Specification-based monitoring of cyber-physical systems: a survey on theory, tools and applications, *Lect. Runtime Verif.: Introd. Adv. Top.* (2018) 135–175.
- [5] R. Yan, A. Julius, Distributed consensus-based online monitoring of robot swarms with temporal logic specifications, *IEEE Robot. Autom. Lett.* 7 (4) (2022) 9413–9420.
- [6] E. Bonnah, K.A. Hoque, Runtime monitoring of time window temporal logic, *IEEE Robot. Autom. Lett.* 7 (3) (2022) 5888–5895.
- [7] C. Eisner, D. Fisman, J. Havlicek, Y. Lustig, A. McIsaac, D. Van Campenhout, Reasoning with temporal logic on truncated paths, in: *International Conference on Computer Aided Verification*, 2003, pp. 27–39.
- [8] A. Dokhanchi, B. Hoxha, G. Fainekos, On-line monitoring for temporal logic robustness, in: *International Conference on Runtime Verification*, 2014, pp. 231–246.
- [9] A. Donzé, T. Ferrere, O. Maler, Efficient robust monitoring for STL, in: *International Conference on Computer Aided Verification*, 2013, pp. 264–279.
- [10] J.V. Deshmukh, A. Donzé, S. Ghosh, X. Jin, G. Juniwal, S.A. Seshia, Robust online monitoring of signal temporal logic, *Form. Methods Syst. Des.* 51 (2017) 5–30.
- [11] Z. Zhang, J. An, P. Arcaini, I. Hasuo, Online causation monitoring of signal temporal logic, in: *International Conference on Computer Aided Verification*, 2023, pp. 62–84.
- [12] G.E. Fainekos, G.J. Pappas, Robustness of temporal logic specifications for continuous-time signals, *Theoret. Comput. Sci.* 410 (42) (2009) 4262–4291.
- [13] H. Abbas, B. Bonakdarpour, Leveraging system dynamics in runtime verification of cyber-physical systems, in: *International Symposium on Leveraging Applications of Formal Methods*, Springer, 2022, pp. 264–278.
- [14] X. Qin, J.V. Deshmukh, Clairvoyant monitoring for signal temporal logic, in: *International Conference Formal Modeling and Analysis of Timed Systems*, Springer, 2020, pp. 178–195.
- [15] M. Ma, J. Stankovic, E. Bartocci, L. Feng, Predictive monitoring with logic-calibrated uncertainty for cyber-physical systems, *ACM Trans. Embed. Comput. Syst.* 20 (5s) (2021) 1–25.
- [16] H. Yoon, S. Sankaranarayanan, Predictive runtime monitoring for mobile robots using logic-based bayesian intent inference, in: *IEEE International Conference on Robotics and Automation*, IEEE, 2021, pp. 8565–8571.
- [17] A. Momtaz, N. Basnet, H. Abbas, B. Bonakdarpour, Predicate monitoring in distributed cyber-physical systems, *Int. J. Softw. Tools Technol. Transf.* 25 (4) (2023) 541–556.
- [18] F. Cairoli, L. Bortolussi, N. Paoletti, Learning-based approaches to predictive monitoring with conformal statistical guarantees, in: *International Conference on Runtime Verification*, Springer, 2023, pp. 461–487.
- [19] L. Lindemann, X. Qin, J.V. Deshmukh, G.J. Pappas, Conformal prediction for stl runtime verification, in: *ACM/IEEE International Conference on Cyber-Physical Systems*, 2023, pp. 142–153.
- [20] Y. Zhao, B. Hoxha, G. Fainekos, J.V. Deshmukh, L. Lindemann, Robust conformal prediction for STL runtime verification under distribution shift, in: *ACM/IEEE International Conference on Cyber-Physical Systems*, 2024, pp. 169–179.
- [21] X. Yu, W. Dong, S. Li, X. Yin, Model predictive monitoring of dynamical systems for signal temporal logic specifications, *Automatica* 160 (2024) 111445.
- [22] C. Wang, X. Yu, J. Zhao, L. Lindemann, X. Yin, Sleep when everything looks fine: Self-triggered monitoring for signal temporal logic tasks, *IEEE Robot. Autom. Lett.* (2024).
- [23] Q.H. Ho, R.B. Ilyes, Z.N. Sunberg, M. Lahijanian, Automaton-guided control synthesis for signal temporal logic specifications, in: *IEEE Conference on Decision and Control*, IEEE, 2022, pp. 3243–3249.
- [24] P. Yu, X. Tan, D. Dimarogonas, Continuous-time control synthesis under nested signal temporal logic specifications, *IEEE Trans. Robot.* 40 (2024) 2272–2286.
- [25] T. Claudet, D. Martire, D. Losa, F. Sanfedino, D. Alazard, A novel graph-based theory for convexification of mission-planning constraints and generative pre-trained trajectory optimization, *IFAC J. Syst. Control.* 30 (2024) 100286.
- [26] O. Maler, D. Nickovic, Monitoring temporal properties of continuous signals, in: *FTRTF*, Springer, 2004, pp. 152–166.
- [27] P. Jagtap, S. Soudjani, M. Zamani, Formal synthesis of stochastic systems via control barrier certificates, *IEEE Trans. Autom. Control* 66 (7) (2020) 3097–3110.