

Online Monitoring of Metric Temporal Logic

Hsi-Ming Ho, Joël Ouaknine and James Worrell

Department of Computer Science, University of Oxford
Wolfson Building, Parks Road, Oxford, OX1 3QD, UK

Abstract. Current approaches to monitoring real-time properties suffer either from unbounded space requirements or lack of expressiveness. In this paper, we adapt a separation technique enabling us to rewrite arbitrary MTL formulas into LTL formulas over a set of atoms comprising bounded MTL formulas. As a result, we obtain the first trace-length independent online monitoring procedure for full MTL.

1 Introduction

In recent years, there has been increasing interest in *runtime verification* as a complement to traditional model checking techniques (see [18, 26] for surveys). Runtime monitoring, for example, may be used in situations in which we wish to evaluate a system that is either too complex to model or whose internal details are not accessible. Moreover, logics whose model-checking problems are undecidable may become tractable in this more restricted setting. The latter is the case in the present paper, which is concerned with runtime monitoring of *Metric Temporal Logic* with both forwards and backwards temporal modalities (MTL[U, S]).

Given an MTL[U, S] formula φ and a finite timed word ρ , the *prefix* problem asks whether all infinite timed words extending ρ satisfy φ . The *monitoring* problem can be seen as an *online* version of the prefix problem where ρ is given incrementally, one event at a time. The monitoring procedure is required to output an answer when either (i) all infinite extensions of the current trace satisfy the specification, or (ii) no infinite extension of the current trace can possibly meet the specification. In this paper, we consider a restricted version of the monitoring problem, based on the notion of *informative prefixes* [17].

Ideally, for a monitoring procedure to be practical, we require that it be *trace-length independent* [5] in the sense that the total space requirement should not depend on the length of the input trace. With this objective in mind, the principal difficulty in monitoring MTL[U, S] is that it allows unbounded intervals and nesting of future and past operators, and hence the truth value of a formula at some point may depend on the truth values of its subformulas arbitrarily far in the future or past. For this reason, most real-time monitoring procedures in the literature consider only certain fragments of MTL[U, S], e.g., only allowing bounded future modalities.¹ An exception is [4] which handles the

¹ Note in passing that, unlike for LTL, past modalities strictly increases the expressiveness of MTL [7].

full logic $\text{MTL}[\mathbf{U}, \mathbf{S}]$, but which unfortunately fails to be trace-length independent. The main contribution of this paper is a novel online monitoring procedure for $\text{MTL}[\mathbf{U}, \mathbf{S}]$. The procedure we give handles the full logic $\text{MTL}[\mathbf{U}, \mathbf{S}]$ and is trace-length independent,² making it suitable for behaviours with potentially unbounded lengths, e.g., network activity logs. For a given formula, we first adapt a separation theorem of [15] to rewrite an $\text{MTL}[\mathbf{U}, \mathbf{S}]$ formula into an $\text{LTL}[\mathbf{U}, \mathbf{S}]$ formula over a set of atoms comprising bounded $\text{MTL}[\mathbf{U}, \mathbf{S}]$ formulas, whose truth values are computed and stored efficiently. The remaining untimed component is then handled via translation to deterministic finite automata. The resulting algorithm is free of dynamic memory allocations, linked lists, etc., and hence can be implemented efficiently. We can moreover leverage efficient techniques for monitoring $\text{LTL}[\mathbf{U}, \mathbf{S}]$, including highly optimised automaton translations via two-way alternating automata [8, 12]. Since we work in the *pointwise semantics*, the techniques in [15] (developed for the *continuous semantics*) must be carefully adapted. Indeed, the main result of [15]—expressive completeness of $\text{MTL}[\mathbf{U}, \mathbf{S}]$ for first-order logic—is known not to hold in the pointwise setting.

2 Related Work

The most closely related work to the present paper is that of Finkbeiner and Kutz [11], which concerns monitoring MTL over a discrete-time semantics. They handle bounded formulas in a similar fashion to us and highlight the problematic role of unbounded temporal operators. However they do not exploit a syntactic rewriting of unbounded operators from the scope of bounded operators, and are forced to apply specialised constructions in this case.

Another highly relevant work is that of Nickovic and Piterman [23], in which a translation from MTL to deterministic timed automata is proposed. The essence of the method is the observation that, while the truth values of unbounded subformulas must necessarily be guessed, the truth values of bounded subformulas can be obtained via bounded look-ahead. In spirit, this is very similar to our approach. The main differences are that they consider only the future fragment, and we handle bounded subformulas explicitly rather than encoding them into clock constraints.

Regarding real-time logics with past, it is known that the non-punctual fragment of $\text{MTL}[\mathbf{U}, \mathbf{S}]$, called $\text{MITL}[\mathbf{U}, \mathbf{S}]$, can be translated into timed automata. Besides the standard constructions [1, 2], there have been some proposals for simplified or improved constructions [9, 16, 20]. The difficulty in using these constructions for monitoring lies in the fact that timed automata cannot be determined in general. In principle one can carry out determinisation on-the-fly for

² As shown in [19], trace-length independence necessarily requires a global bound on the *variability* of time sequences, i.e., the maximum number of events which can occur in any given unit-duration time interval. This is a standard assumption which is in practice always met by physical systems. The proof in [19] is carried out in the continuous semantics, but it goes through in the pointwise case as well.

timed words of bounded variability; however, it is not clear that this approach can yield an efficient procedure.

Automata-free monitoring procedures also appear in the literature. For example, Baldor and Niu propose a monitoring procedure for $\text{MTL}[\mathbf{U}, \mathbf{S}]$ involving a transition-based representation of flows [4]. The main drawback of this procedure, as noted above, and in the paper itself, is that the procedure is not trace-length independent.

3 Background

3.1 Metric Temporal Logic

A *time sequence* $\tau = \tau_1\tau_2\dots$ is a non-empty strictly increasing sequence of rational numbers such that $\tau_1 = 0$. We consider both finite and infinite time sequences, denoting by $|\tau|$ the length of such a sequence. If τ is infinite we require it to be unbounded, i.e., we disallow so-called Zeno sequences.

A *timed word* over a finite alphabet Σ is a pair $\rho = (\sigma, \tau)$, where $\sigma = \sigma_1\sigma_2\dots$ is a non-empty finite or infinite word over Σ and τ is a time sequence of the same length. For $\rho = (\sigma, \tau)$, we refer to the pair (σ_i, τ_i) as the i^{th} event in ρ , and define the *distance* between i^{th} and j^{th} events to be $\tau_j - \tau_i$. We equivalently consider a timed word as a sequence of timed events $(\sigma_1, \tau_1)(\sigma_2, \tau_2)\dots$. The finite timed words considered in this paper arise as prefixes of infinite timed words, and so we sometimes use the word *prefix* to denote an arbitrary finite timed word. We write $T\Sigma^*$ and $T\Sigma^\omega$ for the respective sets of finite and infinite timed words over Σ . For a set of propositions P we write $\Sigma_P = 2^P$.

For a space-bounded online monitoring procedure to be possible, we must impose a global bound on the variability of time sequences, cf. [19]. Henceforth we assume that for some absolute constant k_{var} all timed words have variability at most k_{var} , i.e., there are at most k_{var} events in any unit time interval.

We specify properties of timed words using *Metric Temporal Logic* with both the ‘Until’ and ‘Since’ modalities, denoted $\text{MTL}[\mathbf{U}, \mathbf{S}]$. Given a set of propositions P , the formulas of $\text{MTL}[\mathbf{U}, \mathbf{S}]$ are given by the following grammar

$$\varphi ::= p \mid \mathbf{true} \mid \varphi_1 \wedge \varphi_2 \mid \neg\varphi \mid \varphi_1 \mathbf{U}_I \varphi_2 \mid \varphi_1 \mathbf{S}_I \varphi_2$$

where $p \in P$ and $I \subseteq (0, \infty)$ is an interval with endpoints in $\mathbb{Q}_{\geq 0} \cup \{\infty\}$. We sometime omit the subscript I if $I = (0, \infty)$. Given $x \in \mathbb{Q}$, we write $x < I$ to mean $x < \sup(I)$.

Additional temporal operators and dual operators are defined in the standard way, e.g., $\mathbf{P}_I\varphi \equiv \mathbf{true} \mathbf{S}_I \varphi$ and $\mathbf{H}_I\varphi \equiv \neg\mathbf{P}_I\neg\varphi$. For an $\text{MTL}[\mathbf{U}, \mathbf{S}]$ formula φ , we denote by $|\varphi|$ the number of subformulas of φ .

The satisfaction relation $\rho, i \models \varphi$ for an $\text{MTL}[\mathbf{U}, \mathbf{S}]$ formula φ , an infinite timed word $\rho = (\sigma, \tau)$ and a position $i \geq 1$ is defined as follows:

- $\rho, i \models p$ iff $p \in \sigma_i$
- $\rho, i \models \varphi_1 \mathbf{U}_I \varphi_2$ iff there exists $j > i$ such that $\rho, j \models \varphi_2$, $\tau_j - \tau_i \in I$, and $\rho, k \models \varphi_1$ for all k with $i < k < j$

- $\rho, i \models \varphi_1 \mathbf{S}_I \varphi_2$ iff there exists $j, 1 \leq j < i$ such that $\rho, j \models \varphi_2, \tau_i - \tau_j \in I$ and $\rho, k \models \varphi_1$ for all k with $j < k < i$.³

The semantics of the Boolean connectives is defined in the expected way.

We say that ρ satisfies φ , denoted $\rho \models \varphi$, if $\rho, 1 \models \varphi$. We write $\mathcal{L}(\varphi)$ for the set of infinite timed words that satisfy φ . Abusing notation, we also write $\mathcal{L}(\psi)$ for the set of infinite (untimed) words that satisfy the LTL[\mathbf{U}, \mathbf{S}] formula ψ , and $\mathcal{L}(\mathcal{A})$ for the set of infinite words accepted by \mathcal{A} .

3.2 New Modalities

In this section we recall from [14] the definition of the ‘‘Earlier’’ modalities \mathfrak{E}_I^c and their past counterparts, the ‘‘Newer’’ modalities \mathfrak{N}_I^c . In [14] these were introduced to obtain an extension of MTL[\mathbf{U}, \mathbf{S}] that is expressively complete over bounded timed words for FO[$<, +\mathbb{Q}$], i.e., monadic first-order logic with a binary order relation $<$ and binary distance- q relations for each $q \in \mathbb{Q}$.

Given $c \in \mathbb{N}$ and an interval $I \subseteq (0, \infty)$ with endpoints in $\mathbb{Q}_{\geq 0} \cup \{\infty\}$, define $\mathfrak{E}_I^c \varphi$ to hold at position i of an infinite timed word ρ if there exists j such that $\tau_j - \tau_i \in I, j$ is least such that $\tau_j - \tau_i \in I, \rho, j \models \varphi$, and there is no event with timestamp in $(\tau_i + c, \tau_i + c + (\tau_j - a))$, where a is the left endpoint of I . Intuitively, $\mathfrak{E}_{(a,b)}^c P$ is equivalent to $(\mathbf{F}_{[a,a]} P_\epsilon \wedge \mathbf{F}_{[c,c]} P_\epsilon) \mathbf{U}_{(0,b-a)} (\mathbf{F}_{[a,a]} P)$ in the continuous semantics where P_ϵ is a special proposition that holds at all ‘no-event’ points.

The definition of the past counterpart \mathfrak{N}_I^c is symmetric and omitted; see [14] for details.

We refer to the extended logic as MTL[$\mathbf{U}, \mathbf{S}, \mathfrak{E}, \mathfrak{N}$]. Notations defined in the previous subsection, e.g., $|\varphi|$, extend naturally to MTL[$\mathbf{U}, \mathbf{S}, \mathfrak{E}, \mathfrak{N}$] formulas.

It was shown in [14] that the above two modalities strictly extend the expressiveness of MTL[\mathbf{U}, \mathbf{S}] over bounded timed words, yielding a logic that is expressively complete for FO[$<, +\mathbb{Q}$] in this setting. Here we similarly note that the new modalities cannot be expressed in MTL[\mathbf{U}, \mathbf{S}] over infinite timed words.

Proposition 1. *Over infinite timed words, the modality $\mathfrak{E}_{(1,2)}^1$ cannot be expressed in MTL[\mathbf{U}, \mathbf{S}] with rational constants.*

The proof of Proposition 1 is given in Appendix A.

Finally we remark that \mathfrak{E}_I^c and \mathfrak{N}_I^c are both definable in FO[$<, +\mathbb{Q}$] over infinite timed words. We discuss the question of expressive completeness over infinite (unbounded) timed words in the Conclusion.

3.3 Truncated Semantics and Informative Prefixes

Since in online monitoring one naturally deals with truncated paths, it is useful to define a satisfaction relation of formulas over finite timed words. To this

³ Note that we adopt strict-future and strict-past interpretations to \mathbf{U}_I and \mathbf{S}_I . It is easy to see that, e.g., weak-future until operators can be defined in strict-future ones.

end we adopt a timed version of the *truncated semantics* [10] which incorporates *strong* and *weak* views on satisfaction over truncated paths. These views indicate whether the evaluation of the formula ‘has completed’ on the finite path, i.e., whether the truth value of the formula on the whole path is already determined. For example, the formula $\mathbf{F}_{(0,5)}p$ is weakly satisfied by any finite timed word whose time points are all strictly less than 5 since there is an extension that satisfies the formula. We also consider the *neutral* view, which extends to MTL the traditional LTL semantics over finite words [21].

The respective strong, neutral and weak satisfaction relations will be denoted by \models_f^+ , \models_f and \models_f^- respectively. The definitions below closely follow [10].

Definition 1. *The satisfaction relation $\rho, i \models_f^+ \varphi$ for an MTL[U, S, \mathfrak{E} , \mathfrak{N}] formula φ , a finite timed word $\rho = (\sigma, \tau)$ and a position i , $1 \leq i \leq |\rho|$ is defined as follows:*

- $\rho, i \models_f^+ p$ iff $p \in \sigma_i$
- $\rho, i \models_f^+ \mathbf{true}$
- $\rho, i \models_f^+ \varphi_1 \wedge \varphi_2$ iff $\rho, i \models_f^+ \varphi_1$ and $\rho, i \models_f^+ \varphi_2$
- $\rho, i \models_f^+ \neg\varphi$ iff $(\rho, i) \not\models_f^+ \varphi$
- $\rho, i \models_f^+ \varphi_1 \mathbf{U}_I \varphi_2$ iff there exists j , $i < j \leq |\rho|$ such that $\rho, j \models_f^+ \varphi_2$, $\tau_j - \tau_i \in I$, and $\rho, k \models_f^+ \varphi_1$ for all k with $i < k < j$
- $\rho, i \models_f^+ \varphi_1 \mathbf{S}_I \varphi_2$ iff there exists j , $1 \leq j < i$ such that $\rho, j \models_f^+ \varphi_2$, $\tau_i - \tau_j \in I$ and $\rho, k \models_f^+ \varphi_1$ for all k with $j < k < i$
- $\rho, i \models_f^+ \mathfrak{E}_I^c \varphi$ iff there exists j , $i < j \leq |\rho|$ such that $\rho, j \models_f^+ \varphi$, $\tau_j - \tau_i \in I$, $\tau_{j-1} - \tau_i \notin I$, and there is no k , $i < k < j$ such that $\tau_k - \tau_i > c$ and $\tau_j - \tau_k > (a - c)$, where a is the left endpoint of I
- $\rho, i \models_f^+ \mathfrak{N}_I^c \varphi$ iff there exists j , $1 \leq j < i$ such that $\rho, j \models_f^+ \varphi$, $\tau_i - \tau_j \in I$, $\tau_i - \tau_{j+1} \notin I$, and there is no k , $j < k < i$ such that $\tau_i - \tau_k > c$ and $\tau_k - \tau_j > a - c$, where a is the left endpoint of I .

Definition 2. *The satisfaction relation $\rho, i \models_f^- \varphi$ for an MTL[U, S, \mathfrak{E} , \mathfrak{N}] formula φ , a finite timed word $\rho = (\sigma, \tau)$ and a position i , $1 \leq i \leq |\rho|$ is defined as follows:*

- $\rho, i \models_f^- p$ iff $p \in \sigma_i$
- $\rho, i \models_f^- \mathbf{true}$
- $\rho, i \models_f^- \varphi_1 \wedge \varphi_2$ iff $\rho, i \models_f^- \varphi_1$ and $\rho, i \models_f^- \varphi_2$
- $\rho, i \models_f^- \neg\varphi$ iff $(\rho, i) \not\models_f^- \varphi$
- $\rho, i \models_f^- \varphi_1 \mathbf{U}_I \varphi_2$ iff either of the following holds:
 - there exists j , $i < j \leq |\rho|$ such that $\rho, j \models_f^- \varphi_2$, $\tau_j - \tau_i \in I$, and $\rho, k \models_f^- \varphi_1$ for all k with $i < k < j$
 - $\tau_{|\rho|} - \tau_i < I$ and $\rho, k \models_f^- \varphi_1$ for all k with $i < k \leq |\rho|$
- $\rho, i \models_f^- \varphi_1 \mathbf{S}_I \varphi_2$ iff there exists j , $1 \leq j < i$ such that $\rho, j \models_f^- \varphi_2$, $\tau_i - \tau_j \in I$ and $\rho, k \models_f^- \varphi_1$ for all k with $j < k < i$
- $\rho, i \models_f^- \mathfrak{E}_I^c \varphi$ iff either of the following holds:

- there exists $j, i < j \leq |\rho|$ such that $\rho, j \models_{\bar{f}} \varphi, \tau_j - \tau_i \in I, \tau_{j-1} - \tau_i \notin I$ and there is no $j', i < j' < j$ such that $\tau_{j'} - \tau_i > c$ and $\tau_j - \tau_{j'} > (a - c)$, where a is the left endpoint of I
 - $\tau_{|\rho|} - \tau_i < I$ and for all $j, i < j \leq |\rho|, \tau_j - \tau_i \notin I$
- $(\rho, i) \models_{\bar{f}} \mathfrak{N}_I^c \varphi$ iff there exists $j, 1 \leq j < i$ such that $\rho, j \models_{\bar{f}} \varphi, \tau_i - \tau_j \in I, \tau_i - \tau_{j+1} \notin I$ and there is no $j', j < j' < i$ such that $\tau_i - \tau_{j'} > c$ and $\tau_{j'} - \tau_j > (a - c)$, where a is the left endpoint of I .

The following proposition which helps explain the terms strong, neutral and weak, can be proved by a simple induction on the structure of φ .

Proposition 2. For a finite timed word ρ , a position i in ρ and a formula φ ,

$$\rho, i \models_{\bar{f}}^+ \varphi \rightarrow \rho, i \models_{\bar{f}} \varphi \text{ and } \rho, i \models_{\bar{f}} \varphi \rightarrow \rho, i \models_{\bar{f}}^- \varphi.$$

A closely related notion, *informative prefixes* [17], has been adopted in several works on online monitoring of untimed properties, e.g., [3, 13]. Intuitively, an informative prefix for a formula φ is a prefix that ‘tells the whole story’ about the fulfilment or violation of φ .⁴ We give two examples before the formal definition.

Example 1. Consider the following formula over $\{p_1\}$:

$$\varphi = \mathbf{FG}(\neg p_1) \wedge \mathbf{G}(p_1 \rightarrow \mathbf{F}_{(0,3)} p_1).$$

The finite timed word $\rho = (\{p_1\}, 0)(\{p_1\}, 2)(\emptyset, 5.5)$ is an informative bad prefix for φ , since no extension satisfies the second conjunct. On the other hand, while $\rho' = (\{p_1\}, 0)(\{p_1\}, 2)(\{p_1\}, 4)$ is a bad prefix for φ , it has (different) extensions that satisfy the left and right conjuncts. Thus we do not consider it an informative bad prefix.

Example 2. Consider the following formula over $\{p_1\}$:

$$\varphi' = \mathbf{G}(\neg p_1) \wedge \mathbf{G}(p_1 \rightarrow \mathbf{F}_{(0,3)} p_1).$$

This formula is equivalent to the formula φ in the previous example. However, all bad prefixes for φ' are informative.

The formal definitions are as follows. If a prefix ρ strongly satisfies φ then we say that it is an *informative good prefix* for φ . Similarly we say ρ is an *informative bad prefix* for φ when it fails to weakly satisfy φ . Finally ρ is an *informative prefix* if it is either an informative good prefix or an informative bad prefix. Here we have adopted the semantic characterisation of informative prefixes in terms of the truncated semantics from [10], rather than the original syntactic definition [17].

The following proposition follows immediately from the definition of informative prefixes.

⁴ Our usage of the term *informative* slightly deviates from [17] as in that paper the term refers exclusively to bad prefixes.

Proposition 3. ρ is informative for φ iff ρ is informative for $\neg\varphi$.

Since $\rho \models_f \varphi \leftrightarrow \rho \not\models_f \neg\varphi$, negating a formula essentially exchanges its set of informative good prefixes and informative bad prefixes. The following proposition says ‘something good remains good’ and ‘something bad remains bad’.

Proposition 4. For a finite timed word ρ , a position i in ρ and a formula φ , if ρ is a prefix of the finite timed word ρ' , then

$$\rho, i \models_f^+ \varphi \rightarrow \rho', i \models_f^+ \varphi \text{ and } \rho, i \not\models_f^+ \varphi \rightarrow \rho', i \not\models_f^+ \varphi.$$

4 LTL[U, S] over Bounded Atoms

In this section we present a series of logical equivalences that can be used to rewrite a given MTL[U, S] formula into an equivalent formula in which no unbounded temporal operator occurs within the scope of a bounded operator. Only the rules for future modalities and open intervals are given, as the rules for past modalities are symmetric and the rules for other types of intervals are straightforward variants.

4.1 Normal Form

We say a formula is in *normal form* if it satisfies the following.

- (i) All occurrences of unbounded temporal operators are of the form $\mathbf{U}_{(0,\infty)}$, $\mathbf{S}_{(0,\infty)}$, $\mathbf{G}_{(0,\infty)}$, $\mathbf{H}_{(0,\infty)}$.
- (ii) All other occurrences of temporal operators are of the form \mathfrak{E}_I^c , \mathfrak{N}_I^c , \mathbf{U}_I , \mathbf{S}_I with bounded I .
- (iii) Negation is only applied to propositions or bounded temporal operators (except $\mathbf{G}_{(0,\infty)}$, $\mathbf{H}_{(0,\infty)}$).
- (iv) In any subformula of the form $\varphi_1 \mathbf{U}_I \varphi_2$, $\varphi_1 \mathbf{S}_I \varphi_2$, $\mathbf{F}_I \varphi_2$, $\mathbf{P}_I \varphi_2$ where I is bounded, φ_1 is a disjunction of temporal subformulas and propositions and φ_2 is a conjunction thereof.
- (v) In any subformula of the form $\mathfrak{E}_I^c \varphi$, $\mathfrak{N}_I^c \varphi$, the outermost connective of φ can only be negation or temporal operators.

We describe how to rewrite a given formula into normal form. To satisfy (i) and (ii), apply the usual rules (e.g., $\mathbf{G}_I \varphi \leftrightarrow \neg \mathbf{F}_I \neg \varphi$) and the rule:

$$\varphi_1 \mathbf{U}_{(a,\infty)} \varphi_2 \leftrightarrow \varphi_1 \mathbf{U} \varphi_2 \wedge (\mathbf{F}_{(0,a]} \mathbf{true} \rightarrow \mathbf{G}_{(0,a]} (\varphi_1 \wedge \varphi_1 \mathbf{U} \varphi_2)).$$

Also use the rule:

$$\mathfrak{E}_{(a,\infty)}^c \varphi \leftrightarrow \mathfrak{E}_{(a,a+\max(c,a-c))}^c \varphi \vee (\neg \mathbf{F}_{(c,a]} \mathbf{true} \wedge \mathbf{F}_{[0,c]}^w (\mathbf{false} \mathbf{U}_{(c,\infty)} \varphi)).$$

To satisfy (iii), use the usual rules and the rule:

$$\neg(\varphi_1 \mathbf{U} \varphi_2) \leftrightarrow \mathbf{G} \neg \varphi_2 \vee (\neg \varphi_2 \mathbf{U} (\neg \varphi_2 \wedge \neg \varphi_1)).$$

For (iv) and (v), use the usual rules of Boolean algebra and the rules below:

$$\begin{aligned}
\phi \mathbf{U}_I (\varphi_1 \vee \varphi_2) &\leftrightarrow (\phi \mathbf{U}_I \varphi_1) \vee (\phi \mathbf{U}_I \varphi_2) \\
(\varphi_1 \wedge \varphi_2) \mathbf{U}_I \phi &\leftrightarrow (\varphi_1 \mathbf{U}_I \phi) \wedge (\varphi_2 \mathbf{U}_I \phi) \\
\mathfrak{E}_I^c(\varphi_1 \vee \varphi_2) &\leftrightarrow \mathfrak{E}_I^c \varphi_1 \vee \mathfrak{E}_I^c \varphi_2 \\
\mathfrak{E}_I^c(\varphi_1 \wedge \varphi_2) &\leftrightarrow \mathfrak{E}_I^c \varphi_1 \wedge \mathfrak{E}_I^c \varphi_2.
\end{aligned}$$

4.2 Extracting Unbounded Operators from Bounded Operators

We now provide a set of rewriting rules that extract unbounded operators from the scopes of bounded operators. In what follows, let $\varphi_{xlb} = \mathbf{false} \mathbf{U}_{(0,b)} \mathbf{true}$, $\varphi_{y lb} = \mathbf{false} \mathbf{S}_{(0,b)} \mathbf{true}$ and

$$\begin{aligned}
\varphi_{ugb} &= \left((\varphi_{xlb} \rightarrow \mathbf{G}_{(b,2b)} \varphi_1) \wedge (\neg \varphi_{y lb} \rightarrow (\varphi_1 \wedge \mathbf{G}_{(0,b]} \varphi_1)) \right) \\
&\quad \mathbf{U} \left((\varphi_1 \wedge (\varphi_1 \mathbf{U}_{(b,2b)} \varphi_2)) \vee \left(\neg \varphi_{y lb} \wedge \left(\varphi_2 \vee (\varphi_1 \wedge (\varphi_1 \mathbf{U}_{(0,b]} \varphi_2)) \right) \right) \right), \\
\varphi_{ggb} &= \mathbf{G} \left((\varphi_{xlb} \rightarrow \mathbf{G}_{(b,2b)} \varphi_1) \wedge (\neg \varphi_{y lb} \rightarrow (\varphi_1 \wedge \mathbf{G}_{(0,b]} \varphi_1)) \right).
\end{aligned}$$

Proposition 5. *The following equivalences hold over infinite timed words.*

$$\begin{aligned}
\mathfrak{E}_{(a,b)}^c(\varphi_1 \mathbf{U} \varphi_2) &\leftrightarrow (\mathfrak{E}_{(a,b)}^c \mathbf{true}) \\
&\quad \wedge \left(\mathfrak{E}_{(a,b)}^a(\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2) \vee (\mathfrak{E}_{(a,b)}^a(\mathbf{G}_{(0,2b)} \varphi_1) \wedge \varphi_{ugb}) \right) \\
\mathfrak{E}_{(a,b)}^c \mathbf{G} \varphi &\leftrightarrow (\mathfrak{E}_{(a,b)}^c \mathbf{true}) \wedge (\mathfrak{E}_{(a,b)}^a \mathbf{G}_{(0,2b)} \varphi) \wedge \varphi_{ggb} \\
\mathfrak{E}_{(a,b)}^c(\varphi_1 \mathbf{S} \varphi_2) &\leftrightarrow (\mathfrak{E}_{(a,b)}^c \mathbf{true}) \\
&\quad \wedge \left((\mathfrak{E}_{(a,b)}^a(\varphi_1 \mathbf{S}_{(0,b)} \varphi_2)) \vee ((\mathfrak{E}_{(a,b)}^a \mathbf{H}_{(0,b)} \varphi_1) \wedge (\varphi_1 \mathbf{S} \varphi_2)) \right) \\
\mathfrak{E}_{(a,b)}^c \mathbf{H} \varphi &\leftrightarrow (\mathfrak{E}_{(a,b)}^c \mathbf{true}) \wedge (\mathfrak{E}_{(a,b)}^a \mathbf{H}_{(0,b)} \varphi) \wedge \mathbf{H} \varphi.
\end{aligned}$$

Proof. We prove the first rule as the proofs of other rules are similar. In the following, let the current position be i and the position of the first event in $(\tau_i + a, \tau_i + b)$ be j .

For the forward direction, let the witness position where φ_2 holds be w . If $\tau_w < \tau_j + 2b$, the subformula $\mathfrak{E}_{(a,b)}^a(\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2)$ clearly holds at i . Otherwise we have $\mathfrak{E}_{(a,b)}^a(\mathbf{G}_{(0,2b)} \varphi_1)$ holds at i . It follows that $(\varphi_{xlb} \rightarrow \mathbf{G}_{(b,2b)} \varphi_1)$ and $\varphi_{y lb}$ (and hence $\neg \varphi_{y lb} \rightarrow (\varphi_1 \wedge \mathbf{G}_{(0,b]} \varphi_1)$) holds at all positions j' , $i < j' < j$. Let $l > j$ the first position such that $\tau_w \in (\tau_l + b, \tau_l + 2b)$. Consider the following cases:

- There is such l : It is clear that $(\varphi_1 \wedge (\varphi_1 \mathbf{U}_{(b,2b)} \varphi_2))$ holds at l . Since $\mathbf{G}_{(b,2b)} \varphi_1$ holds at all positions j'' , $j \leq j'' < l$ by the minimality of l , $(\varphi_{xlb} \rightarrow \mathbf{G}_{(b,2b)} \varphi_1)$ also holds at these positions. For the other conjunct, note that $\varphi_{y lb}$ holds at j and $\varphi_1 \wedge \mathbf{G}_{(0,b]} \varphi_1$ holds at all positions j''' , $j < j''' < l$.

- There is no such l : Consider the following cases:
 - $\neg\varphi_{y|b}$ and $\neg\mathbf{P}_{[b,b]}\mathbf{true}$ holds at w : There is no event in $(\tau_w - 2b, \tau_w)$. The proof is similar to the case where l exists.
 - $\neg\varphi_{y|b}$ and $\mathbf{P}_{[b,b]}\mathbf{true}$ holds at w : Let l' be the position such that $\tau_{l'} = \tau_w - b$. There must be no event in $(\tau_{l'} - b, \tau_{l'})$. It follows that $\neg\varphi_{y|b}$ and $(\varphi_1 \wedge (\varphi_1 \mathbf{U}_{(0,b)} \varphi_2))$ holds at l' . The proof is similar.
 - $\varphi_{y|b}$ holds at w : By assumption, there is no event in $(\tau_w - 2b, \tau_w - b)$. It is easy to see that there is a position such that $\neg\varphi_{y|b} \wedge (\varphi_1 \wedge (\varphi_1 \mathbf{U}_{(0,b)} \varphi_2))$ holds. The proof is again similar.

The other direction is more straightforward. □

Proposition 6. *The following equivalences hold over infinite timed words.*

$$\begin{aligned}
\theta \mathbf{U}_{(a,b)} ((\varphi_1 \mathbf{U} \varphi_2) \wedge \chi) &\leftrightarrow \theta \mathbf{U}_{(a,b)} ((\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2) \wedge \chi) \\
&\quad \vee \left((\theta \mathbf{U}_{(a,b)} (\mathbf{G}_{(0,2b)} \varphi_1 \wedge \chi)) \wedge \varphi_{ugb} \right) \\
\theta \mathbf{U}_{(a,b)} (\mathbf{G}\varphi \wedge \chi) &\leftrightarrow (\theta \mathbf{U}_{(a,b)} (\mathbf{G}_{(0,2b)} \varphi \wedge \chi)) \wedge \varphi_{ggb} \\
\theta \mathbf{U}_{(a,b)} ((\varphi_1 \mathbf{S} \varphi_2) \wedge \chi) &\leftrightarrow \theta \mathbf{U}_{(a,b)} ((\varphi_1 \mathbf{S}_{(0,b)} \varphi_2) \wedge \chi) \\
&\quad \vee \left((\theta \mathbf{U}_{(a,b)} (\mathbf{H}_{(0,b)} \varphi_1 \wedge \chi)) \wedge \varphi_1 \mathbf{S} \varphi_2 \right) \\
\theta \mathbf{U}_{(a,b)} (\mathbf{H}\varphi \wedge \chi) &\leftrightarrow (\theta \mathbf{U}_{(a,b)} (\mathbf{H}_{(0,b)} \varphi \wedge \chi)) \wedge \mathbf{H}\varphi \\
((\varphi_1 \mathbf{U} \varphi_2) \vee \chi) \mathbf{U}_{(a,b)} \theta &\leftrightarrow ((\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2) \vee \chi) \mathbf{U}_{(a,b)} \theta \\
&\quad \vee \left(\left(((\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2) \vee \chi) \mathbf{U}_{(0,b)} (\mathbf{G}_{(0,2b)} \varphi_1) \right) \right. \\
&\quad \quad \wedge \\
&\quad \quad \left. \mathbf{F}_{(a,b)} \theta \wedge \varphi_{ugb} \right) \\
((\mathbf{G}\varphi) \vee \chi) \mathbf{U}_{(a,b)} \theta &\leftrightarrow \chi \mathbf{U}_{(a,b)} \theta \\
&\quad \vee (\chi \mathbf{U}_{(0,b)} (\mathbf{G}_{(0,2b)} \varphi_1) \wedge \mathbf{F}_{(a,b)} \theta \wedge \varphi_{ggb}) \\
((\varphi_1 \mathbf{S} \varphi_2) \vee \chi) \mathbf{U}_{(a,b)} \theta &\leftrightarrow ((\varphi_1 \mathbf{S}_{(0,b)} \varphi_2) \vee \chi) \mathbf{U}_{(a,b)} \theta \\
&\quad \vee \left(\left((\mathbf{H}_{(0,b)} \varphi_1 \vee (\varphi_1 \mathbf{S}_{(0,b)} \varphi_2) \vee \chi) \mathbf{U}_{(a,b)} \theta \right) \right. \\
&\quad \quad \wedge \\
&\quad \quad \left. \varphi_1 \mathbf{S} \varphi_2 \right) \\
((\mathbf{H}\varphi) \vee \chi) \mathbf{U}_{(a,b)} \theta &\leftrightarrow \chi \mathbf{U}_{(a,b)} \theta \vee \left(((\mathbf{H}_{(0,b)} \varphi \vee \chi) \mathbf{U}_{(a,b)} \theta) \wedge \mathbf{H}\varphi \right).
\end{aligned}$$

Proposition 7. *For a formula φ , we can use the rules above to obtain an equivalent formula $\hat{\varphi}$ in which no unbounded temporal operator appears in the scope of a bounded temporal operator.*

Proof. Define the *unbounding depth* $ud(\varphi)$ of a formula φ to be the modal depth of φ counting only unbounded operators. We demonstrate a rewriting process on φ which terminates in an equivalent formula $\hat{\varphi}$ such that any subformula $\hat{\psi}$ of $\hat{\varphi}$ with outermost operator bounded has $ud(\hat{\psi}) = 0$.

Assume that the input formula φ is in normal form. Let k be the largest unbounding depth among all subformulas of φ with bounded outermost operators. We pick all minimal (wrt. subformula order) such subformulas ψ with $ud(\psi) = k$. By applying the rules in Section 4.2, we can rewrite ψ into ψ' where all subformulas of ψ' with bounded outermost operators have unbounding depths strictly less than k . We then substitute these ψ' back into φ to obtain φ' . We repeat this step until there remain no bounded operators with unbounding depth k . Rules that rewrite a formula into normal form are used whenever necessary on relevant subformulas—this will never affect their unbounding depths. It is easy to see that we will eventually obtain such a formula φ^* . Now rewrite φ^* into normal form and start over again. This is to be repeated until we reach $\hat{\varphi}$. \square

Given the input formula φ over propositions $P = \{p_1, \dots, p_n\}$, we can apply the rewriting process above to obtain a formula $\hat{\varphi}$. Since each rewriting rule is a logical equivalence, we have the following theorem.

Theorem 1. $\mathcal{L}(\varphi) = \mathcal{L}(\hat{\varphi})$.

The separation of the original formula could possibly induce a non-elementary blow-up. However, we argue that this rarely happens in practice. For example, consider the following formula:

$$\mathbf{G}(\text{ChangeGear} \rightarrow \mathbf{F}_{(0,30)}(\text{InjectFuel} \wedge \mathbf{P}\text{InjectLubricant})).$$

The syntactically separated version of the formula is

$$\mathbf{G}[\text{ChangeGear} \rightarrow \mathbf{F}_{(0,30)}(\text{InjectFuel} \wedge \mathbf{P}_{(0,30)}\text{InjectLubricant}) \vee (\mathbf{F}_{(0,30)}(\text{InjectFuel}) \wedge \mathbf{P}\text{InjectLubricant})],$$

which is of comparable size.

5 Online Monitoring Procedure

Having obtained $\hat{\varphi} = \Phi(\psi_1, \dots, \psi_m)$ where ψ_1, \dots, ψ_m are bounded formulas over P and Φ is an LTL[U, S] formula, we now introduce new propositions $Q = \{q_1, \dots, q_m\}$ that correspond to bounded subformulas. In this way, we can monitor Φ as an untimed property over Q , only that now we obtain the truth values of q_1, \dots, q_m by simple dynamic programming procedures. As these propositions correspond to bounded formulas, we only need to store a ‘sliding window’ on the input timed word.

5.1 Untimed LTL[U, S] Part

We describe briefly the standard way to construct automata that detect informative prefixes [17]. For a given LTL formula Θ , first use a standard construction [27] to obtain a language-equivalent alternating Büchi automaton \mathcal{A}_Θ . Then redefine its set of accepting states to be the empty set and treat it as an automaton over finite words. The resulting automaton $\mathcal{A}_\Theta^{true}$ accepts exactly all informative good prefixes for Θ . For online monitoring, one can then determinise $\mathcal{A}_\Theta^{true}$ with the usual subset construction. The same can be done for $\neg\Theta$ to obtain a deterministic automaton detecting informative bad prefixes for Θ .

In our case, we first translate the LTL[U, S] formulas Φ and $\neg\Phi$ into a pair of *two-way* alternating Büchi automata. It is easy to see that, with the same ‘tweaks’, we can obtain two automata that accept informative good prefixes and informative bad prefixes for Φ (by Proposition 3). We then apply existing procedures that translate two-way alternating automata over finite words into deterministic automata, e.g., [6]. We call the resulting automata \mathcal{D}_{good} and \mathcal{D}_{bad} and execute them in parallel.

5.2 Bounded Metric Part

We define $fr(\varphi)$ and $pr(\varphi)$ (*future-reach* and *past-reach*) for an MTL[U, S, \mathfrak{E} , \mathfrak{N}] formula φ as follows (the cases for boolean connectives are defined as expected):

- $fr(\mathbf{true}) = pr(\mathbf{true}) = fr(p) = pr(p) = 0$ for all $p \in P$
- $fr(\varphi_1 \mathbf{U}_I \varphi_2) = \sup(I) + \max(fr(\varphi_1), fr(\varphi_2))$
- $pr(\varphi_1 \mathbf{S}_I \varphi_2) = \sup(I) + \max(pr(\varphi_1), pr(\varphi_2))$
- $fr(\varphi_1 \mathbf{S}_I \varphi_2) = \max(fr(\varphi_1), fr(\varphi_2) - \inf(I))$
- $pr(\varphi_1 \mathbf{U}_I \varphi_2) = \max(pr(\varphi_1), pr(\varphi_2) - \inf(I))$
- $fr(\mathfrak{E}_I^c \varphi) = \sup(I) + fr(\varphi)$
- $pr(\mathfrak{N}_I^c \varphi) = \sup(I) + pr(\varphi)$
- $fr(\mathfrak{N}_I^c \varphi) = \max(0, fr(\varphi) - \inf(I))$
- $pr(\mathfrak{E}_I^c \varphi) = \max(0, pr(\varphi) - \inf(I))$.

Intuitively, these indicate the lengths of the time horizons needed to determine the truth value of φ . For a bounded formula ψ , we rewrite ψ so that all temporal operators appear in their positive form, e.g., $\varphi_1 \mathbf{U}_I \varphi_2$. Recall that we assume that timed words are of bounded variability k_{var} . Let $l_f(\psi) = k_{var} \cdot \lceil fr(\psi) \rceil$ and $l_p(\psi) = k_{var} \cdot \lceil pr(\psi) \rceil$.

Naïve Method Suppose that we would like to obtain the truth value of q_i at position j in the input (infinite) timed word $\rho = (\sigma, \tau)$. Observe that only events occurring between $\tau_j - pr(\psi_i)$ and $\tau_j + fr(\psi_i)$ can affect the truth value of ψ_i at j . This implies that $\rho, j \models \psi_i \leftrightarrow \rho', j \models_f \psi_i$, given that ρ' is a prefix of ρ that contains all events between $\tau_j - pr(\psi_i)$ and $\tau_j + fr(\psi_i)$. Since ρ is of bounded variability k_{var} , there will be at most $l_p(\psi) + 1 + l_f(\psi)$ events between $\tau_j - pr(\psi_i)$ and $\tau_j + fr(\psi_i)$. It follows that we can simply record all events

in this interval. Events outside of this interval are irrelevant as they do not affect whether $\rho', j \models_{\mathcal{F}} \psi_i$. In particular, we maintain a two-dimensional array of $l = l_p(\psi) + 1 + l_f(\psi) + 1$ rows and $1 + |\psi|$ columns. The first column is used to store timestamps of the corresponding events.⁵ The last $|\psi|$ columns are used to store the truth values of subformulas. We then use dynamic programming procedures (cf. [22]) to evaluate whether $\rho', j \models_{\mathcal{F}} \psi_i$. These procedures fill up the array in a bottom-up manner, starting from minimal subformulas. New operators $\mathfrak{E}_I^c, \mathfrak{N}_I^c$ can be handled similarly, and the columns for boolean combinations can be filled in the natural way.

Now consider all propositions in Q . We can obtain the truth values of them at all positions in the ‘sliding window’ by using an array of $l_p^Q + 1 + l_f^Q + 1$ rows and $1 + |\psi_1| + \dots + |\psi_m|$ columns, where $l_p^Q = \max_{i \in [1, m]} l_p(\psi_i)$ and $l_f^Q = \max_{i \in [1, m]} l_f(\psi_i)$. Each column can be filled in time linear in its length. Overall, we need an array of size $O(k_{var} \cdot c_{sum} \cdot |\hat{\varphi}|)$ where c_{sum} is the sum of the constants in $\hat{\varphi}$, and for each position j we need time $O(k_{var} \cdot c_{sum} \cdot |\hat{\varphi}|)$ to obtain the truth values of all propositions in Q . This method is not very efficient as for each j we need to fill all columns for temporal subformulas from scratch. Previously computed entries cannot always be reused as certain entries are ‘wrong’—they were computed without the knowledge of events outside of the interval.

Incremental Evaluation We describe an optimisation which allows effective reuse of computed entries stored in the table. The idea is to treat entries that depend on future events as ‘unknown’ and not to fill them. By construction, these unknown entries will not be needed for the result of the evaluation.

For a past subformula, e.g. $\varphi_1 \mathbf{S}_{(a,b)} \varphi_2$, we can simply suspend the column-filling procedure when we filled all entries using the truth values of φ_1 and φ_2 (at various positions) that are currently known. We may continue when the truth values of φ_1 and φ_2 (at some other positions) that are previously unknown become available. Subformulas of the form $\mathfrak{N}_I^c \varphi$ can be handled in a similar way. The case for future subformulas is more involved. Suppose that we filling a column for $p_1 \mathbf{U}_{(a,b)} p_2$ with the naïve method. Denote the corresponding timestamp of an index i in the column by $t(i)$ and the timestamp of the last acquired event by t_{max} . Observe that not all of the truth values at indices $j, t(j) + b > t_{max}$ can be reused later, as they might depend on future events. However, if we know that φ_1 does not hold at some $j', t(j') + b > t_{max}$, then all the truth values at indices $< j'$ can be reused in the following iterations as they cannot depend on future events. Now consider the general case of filling the column for $\psi = \varphi_1 \mathbf{U}_{(a,b)} \varphi_2$. We keep an index j_ψ that points to the first unknown entry in the column, and we now let $t_{max} = \min(t(j_{\varphi_1} - 1), t(j_{\varphi_2} - 1))$. In each iteration, if j_{φ_1} and j_{φ_2} are updated to some new values, t_{max} also changes accordingly. In this case, we first check if $t(j_\psi) + b > t_{max}$. If this is the case, we do nothing (observe the fact that φ_1 must hold at all indices $l, t(j_\psi) < t(l) \leq t_{max}$,

⁵ We assume the timestamps can be finitely represented, e.g., with a built-in data type, and additions and subtractions on them can be done in constant time.

thus the truth value at j_ψ must remain unknown). Otherwise we find the least index $l' > j_\psi$ such that $t(l') + b > t_{max}$. Additionally, we check if all truth values of φ_1 between t_{max} and t_{max}^{old} are **true**, starting from t_{max} . If φ_1 is not satisfied at some position j' then start filling at $\max(l', j') - 1$. Otherwise we start filling from $l' - 1$. Note that we still make use of the truth values of φ_1 and φ_2 up to t_{max} .

Observe that we can use a variable to keep track of the least index $l' > j_\psi$ such that $t(l') + b > t_{max}$ instead of finding it each time since it increases monotonically. Also we can keep track of the greatest index where φ_2 holds. With these variables, we can easily make the extra ‘sweeping’ happen only twice (once for φ_1 and once for φ_2) over newly acquired truth values. Also observe that the truth value of a subformula at a certain position will be filled only once. These observations imply that each entry in the array can be filled in amortised constant time. Similar ideas work for subformulas of the form $\mathfrak{E}_f^c \varphi$. Assuming that each step of a deterministic automaton takes constant time, we can state the following theorem.

Theorem 2. *For a formula φ and an infinite timed word of bounded variability k_{var} , our online monitoring procedure requires two DFAs of size $2^{2^{O(|\Phi|)}}$, an array of size $O(k_{var} \cdot c_{sum} \cdot |\hat{\varphi}|)$ where c_{sum} is the sum of the constants in $\hat{\varphi}$, and amortised time $O(|\hat{\varphi}|)$ per event.*

5.3 Correctness

One may think of the monitoring process on an infinite timed word $\rho \in T\Sigma_P^\omega$ as continuously extending a corresponding finite timed word $\rho' \in T\Sigma_Q^*$. Suppose that, instead of \mathcal{D}_{good} and \mathcal{D}_{bad} , we now execute a deterministic ω -automaton \mathcal{D}_Φ such that $\mathcal{L}(\mathcal{D}_\Phi) = \mathcal{L}(\Phi)$. Since we are implicitly ensuring that the truth values of propositions in Q are valid along the way, it is easy to see that the corresponding run on \mathcal{D}_Φ will be accepting iff $\rho \models \varphi$. However, for the purpose of online monitoring, we will be more interested in deciding whether $\rho \models \varphi$ given only a finite prefix of ρ . In this subsection we show that our approach is both sound and complete for detecting informative prefixes.

The following proposition is immediate since three views of the truncated semantics coincide in this case.

Proposition 8. *For a bounded formula ψ , a finite timed word $\rho = (\sigma, \tau)$ and a position $1 \leq i \leq |\rho|$ such that $\tau_i + fr(\psi) \leq \tau_{|\rho|}$ and $\tau_i - pr(\psi) \geq 0$, we have*

$$\rho, i \models_f^+ \psi \leftrightarrow \rho, i \models_f \psi \leftrightarrow \rho, i \models_f^- \psi.$$

The following lemma implies that the rewriting process outlined in Section 4 preserves the ‘informativeness’ of prefixes.

Lemma 1. *For a formula φ , let φ' be the formula obtained after applying one of the rewriting rules in Section 4 on some of its subformula. We have*

$$\rho \models_f^+ \varphi \leftrightarrow \rho \models_f^+ \varphi' \text{ and } \rho \models_f^- \varphi \leftrightarrow \rho \models_f^- \varphi'.$$

Proof. See Appendix B. □

Given the lemma above, we can state the following theorem.

Theorem 3. *The set of informative good prefixes of φ coincides with the set of informative good prefixes of $\hat{\varphi}$. Same for informative bad prefixes.*

Now we state the main result of the paper in the following two theorems.

Theorem 4 (Soundness). *In our procedure, if we ever reach an accepting state of \mathcal{D}_{good} (\mathcal{D}_{bad}) over a finite word $u \in \Sigma_Q^*$, then the finite timed word $\rho \in T\Sigma_P^*$ that we have read must be an informative good (bad) prefix for φ .*

Proof. For such u and the corresponding ρ (note that $|u| \leq |\rho|$),

$$\forall i \in [1, |u|] \left((u, i \not\models_f^+ \Theta \rightarrow \rho, i \not\models_f^+ \vartheta) \wedge (u, i \models_f^+ \Theta \rightarrow \rho, i \models_f^+ \vartheta) \right)$$

where Θ is a subformula of Φ and $\vartheta = \Theta(\psi_1, \dots, \psi_m)$. This can easily be proved by structural induction. If u is accepted by \mathcal{D}_{good} , we have $u \models_f^+ \Phi$ by construction. By the above we have $\rho \models_f^+ \Phi(\psi_1, \dots, \psi_m)$, as desired. The case for \mathcal{D}_{bad} is symmetric. □

Theorem 5 (Completeness). *Whenever we read an informative good (bad) prefix $\rho = (\sigma, \tau)$ for φ , \mathcal{D}_{good} (\mathcal{D}_{bad}) must eventually reach an accepting state.*

Proof. For the finite word u' obtained a bit later with $|u'| = |\rho|$,

$$\forall i \in [1, |u'|] \left((\rho, i \models_f^+ \vartheta \rightarrow u', i \models_f^+ \Theta) \wedge (\rho, i \not\models_f^+ \vartheta \rightarrow u', i \not\models_f^+ \Theta) \right)$$

where Θ is a subformula of Φ and $\vartheta = \Theta(\psi_1, \dots, \psi_m)$. Again, this can be proved by structural induction (the base step holds by Proposition 4). The theorem follows. □

Remark 1. As pointed out in Example 1, is possible that some of the bad prefixes for φ are not informative. To avoid such a situation, we can impose certain syntactic restrictions on φ . For example, it can be shown that all bad prefixes of Safety-MTL [24] formulas will inevitably be extended to informative bad prefixes.⁶

6 Conclusion

We have given an efficient online monitoring procedure for $\text{MTL}[\mathbf{U}, \mathbf{S}, \mathfrak{E}, \mathfrak{N}]$, the extension of $\text{MTL}[\mathbf{U}, \mathbf{S}]$ with the new modalities. Along the way, we also obtain a new proof that $\text{MTL}[\mathbf{U}, \mathbf{S}]$ is not expressively complete for $\text{FO}[\langle, +\mathbb{Q}]\rangle$ in the pointwise semantics, in contrast to the situation in the continuous semantics. In

⁶ In the words of Kupferman and Vardi [17], all Safety-MTL properties are either *intentionally safe* or *accidentally safe*.

particular, we have introduced two new families of first-order definable modalities $\mathfrak{E}_f^c, \mathfrak{N}_f^c$ that strictly increase the expressiveness of $\text{MTL}[\mathbf{U}, \mathbf{S}]$.

Our procedure detects all informative prefixes for the input formula. It works by rewriting the input $\text{MTL}[\mathbf{U}, \mathbf{S}, \mathfrak{E}, \mathfrak{N}]$ formula into an $\text{LTL}[\mathbf{U}, \mathbf{S}]$ formula over a set of bounded $\text{MTL}[\mathbf{U}, \mathbf{S}, \mathfrak{E}, \mathfrak{N}]$ atoms. The former is converted into a deterministic (untimed) automaton, while the truth values of the latter are maintained by dynamic programming. Crucially we circumvent the potentially delicate issue of translating $\text{MTL}[\mathbf{U}, \mathbf{S}, \mathfrak{E}, \mathfrak{N}]$ to a class of deterministic timed automata.

In the future we plan to compare the expressiveness of $\text{MTL}[\mathbf{U}, \mathbf{S}, \mathfrak{E}, \mathfrak{N}]$ and $\text{FO}[\langle, +\mathbb{Q} \rangle]$. The new modalities $\mathfrak{E}_f^c, \mathfrak{N}_f^c$ seem to fill a certain gap between pointwise and continuous interpretations of $\text{MTL}[\mathbf{U}, \mathbf{S}]$, suggesting that their inclusion may yield expressive completeness. Indeed, we have proved recently that $\text{MTL}[\mathbf{U}, \mathbf{S}, \mathfrak{E}, \mathfrak{N}]$ is expressively complete for $\text{FO}[\langle, +1 \rangle]$ over timed words of bounded durations [14]. Another possible direction for future work is to improve the monitoring procedure. For example, the dynamic programming procedures in Section 5.2 can support subformulas with unbounded past. This can be exploited to use a smaller equivalent formula in place of $\hat{\varphi}$.

References

1. Alur, R., Feder, T., Henzinger, T.: The benefits of relaxing punctuality. *Journal of the ACM* 43(1), 116–146 (1996)
2. Alur, R., Henzinger, T.: Back to the future: towards a theory of timed regular languages. In: *Proceedings of FOCS 1992*. pp. 177–186. IEEE Computer Society Press (1992)
3. Armoni, R., Korchemny, D., Tiemeyer, A., Vardi, M.Y., Zbar, Y.: Deterministic dynamic monitors for linear-time assertions. In: *Proceedings of FATES/RV 2006*. LNCS, vol. 4262, pp. 163–177. Springer (2006)
4. Baldor, K., Niu, J.: Monitoring dense-time, continuous-semantics, metric temporal logic. In: *Proceedings of RV 2012*. LNCS, vol. 7687, pp. 245–259. Springer (2012)
5. Bauer, A., Küster, J., Vegliach, G.: From propositional to first-order monitoring. In: *Proceedings of RV 2013*. LNCS, vol. 8174, pp. 59–75. Springer (2013)
6. Birget, J.C.: State-complexity of finite-state devices, state compressibility and incompressibility. *Mathematical Systems Theory* 26(3), 237–269 (1993)
7. Bouyer, P., Chevalier, F., Markey, N.: On the expressiveness of TPSTL and MTL. In: *Proceedings of FSTTCS 2005*. LNCS, vol. 3821, pp. 432–443. Springer (2005)
8. Dax, C., Klaedtke, F., Lange, M.: On regular temporal logics with past. *Acta Informatica* 47(4), 251–277 (2010)
9. D’Souza, D., Matteplackel, R.: A clock-optimal hierarchical monitoring automaton construction for MITL. Tech. Rep. 2013-1, Indian Institute of Science (2013), <http://www.csa.iisc.ernet.in/TR/2013/1/lics2013-tr.pdf>
10. Eisner, C., Fisman, D., Havlicek, J., Lustig, Y.: Reasoning with temporal logic on truncated paths. In: *Proceedings of CAV 2003*. LNCS, vol. 2725, pp. 27–39. Springer (2003)
11. Finkbeiner, B., Kuhtz, L.: Monitor circuits for LTL with bounded and unbounded future. In: *Proceedings of RV 2009*. LNCS, vol. 5779, pp. 60–75. Springer (2009)

12. Gastin, P., Oddoux, D.: LTL with past and two-way very-weak alternating automata. In: Proceedings of MFCS 2003. LNCS, vol. 2747, pp. 439–448. Springer (2003)
13. Geilen, M.: On the construction of monitors for temporal logic properties. *Electronic Notes in Theoretical Computer Science* 55(2), 181–199 (2001)
14. Ho, H.M., Ouaknine, J.: On the expressiveness of metric temporal logic over bounded timed words (2013), <http://www.cs.ox.ac.uk/people/hsi-ming.ho/exp.pdf>, submitted.
15. Hunter, P., Ouaknine, J., Worrell, J.: Expressive completeness of metric temporal logic. In: Proceedings of LICS 2013. pp. 349–357. IEEE Computer Society Press (2013)
16. Kini, D., Krishna, S., Pandya, P.: On construction of safety signal automata for MITL[U,S] using temporal projections. In: Proceedings of FORMATS 2011. LNCS, vol. 6919, pp. 225–239. Springer (2011)
17. Kupferman, O., Vardi, M.Y.: Model checking of safety properties. *Formal Methods in System Design* 19(3), 291–314 (2001)
18. Leucker, M., Schallhart, C.: A brief account of runtime verification. *Journal of Logic and Algebraic Programming* 78(5), 293–303 (2009)
19. Maler, O., Nickovic, D., Pnueli, A.: Real time temporal logic: Past, present, future. In: Proceedings of FORMATS 2005. LNCS, vol. 3829, pp. 2–16. Springer (2005)
20. Maler, O., Nickovic, D., Pnueli, A.: From MITL to timed automata. In: Proceedings of FORMATS 2006. LNCS, vol. 4202, pp. 274–289. Springer (2006)
21. Manna, Z., Pnueli, A.: *Temporal verification of reactive systems: safety*, vol. 2. Springer (1995)
22. Markey, N., Raskin, J.: Model checking restricted sets of timed paths. *Theoretical Computer Science* 358(2-3), 273–292 (2006)
23. Nickovic, D., Piterman, N.: From MTL to deterministic timed automata. In: Proceedings of FORMATS 2010. LNCS, vol. 6246, pp. 152–167. Springer (2010)
24. Ouaknine, J., Worrell, J.: Safety metric temporal logic is fully decidable. In: Proceedings of TACAS 2006. LNCS, vol. 3920, pp. 411–425. Springer (2006)
25. Pandya, P.K., Shah, S.S.: On expressive powers of timed logics: Comparing boundedness, non-punctuality and deterministic freezing. In: Proceedings of CONCUR 2011. LNCS, vol. 6901, pp. 60–75. Springer (2011)
26. Sokolsky, O., Havelund, K., Lee, I.: Introduction to the special section on runtime verification. *International Journal on Software Tools for Technology Transfer* 14(3), 243–247 (2011)
27. Vardi, M.Y.: An automata-theoretic approach to linear temporal logic. In: *Logics for Concurrency – Structure versus Automata* (8th Banff Higher Order Workshop’95). LNCS, vol. 1043, pp. 238–266. Springer (1996)

A Proof of Proposition 1

In order to prove this proposition, we will need some extra definitions. Given a formula φ , we define its *granularity* g as the largest positive rational such that $\frac{1}{g} \in \mathbb{N}_{>0}$ and all constants appear in φ are integer multiples of g . We then resort to a rational variant of the MTL[U, S] EF theorem (see [25] for details). We write $\rho, i \approx_n^{c \cdot g} \rho', i'$ if for every n -round MTL[U, S] EF-game allowing constants $m \cdot g \leq c \cdot g, m \in \mathbb{N}_{\geq 0}$ on ρ, ρ' that starts from configuration (i, i') , the *Duplicator* always has a winning strategy.

Theorem 6 ([25]). *For infinite timed words ρ, ρ' and an MTL[U, S] formula φ of granularity g , modal depth $\leq n$ with largest constant $\leq c \cdot g$,*

$$\rho, 1 \approx_n^{c \cdot g} \rho', 1 \text{ implies } \rho \models \varphi \leftrightarrow \rho' \models \varphi.$$

We will construct two families of infinite timed words $(M_{n,c,g})$ and $(N_{n,c,g})$ that are indistinguishable by MTL[U, S] formulas of granularity g , modal depth $\leq n$ with all constants $\leq c \cdot g$.

Fix some $n, c \in \mathbb{N}_{\geq 0}$ and $g \in \mathbb{Q}_{>0}$ (wlog. we assume $g < \delta$ and $g < \frac{1}{5}$). Let c' be the least integer greater than $\frac{1-g}{1-2g} \cdot (c+3) + 1$ and pick $\epsilon < \frac{g}{\frac{1}{g}-1}$.

A sequence of ‘empty’ events (in which none of p_1, p_2 hold) starts at time 0 and is equally separated by g . After the last empty event at $c \cdot g$, a number of overlapping segments start at time $(c+1) \cdot g$. Each segment consists of a p_1 -event and a p_2 -event. If the p_1 -event in the i^{th} segment is at, say, t , then its p_2 -event is at $t + 2g + \frac{i}{3 \cdot n \cdot c' + 3} \cdot \epsilon$ (see Figure 1). All p_1 -events in neighbouring segments are separated by $g - \frac{g}{\frac{1}{g}-1}$. We put a total of $2 \cdot n \cdot c' + 1$ segments.⁷

Then an infinite sequence of empty events, equally separated by g , starts at g after the p_2 -event in the last segment. The only difference of the constructed timed words $M_{n,c,g}, N_{n,c,g}$ lies in the middle (i.e., $(n \cdot c' + 1)^{\text{th}}$) segment. Say this segment starts at t , then in $M_{n,c,g}$ we move the corresponding p_2 -event to $t + 2g - \frac{n \cdot c' + 1}{3 \cdot n \cdot c' + 3} \cdot \epsilon$ instead. For convenience, we denote this pair of events in $M_{n,c,g}$ and $N_{n,c,g}$ by x and y and write their corresponding timestamps as t_x and t_y (see Figure 2). It is easy to see that no p_2 -event has distance an integer multiple of g to some other p_1 -event or p_2 -event. This completes our description of the construction of $(M_{n,c,g})$ and $(N_{n,c,g})$.

We say a configuration (i, j) is *identical* if $i = j$. For a position $i \geq 1$, $seg(i)$ indicates to which segment the i^{th} event belongs. For convenience we define $seg(i) = 0$ if the i^{th} event is empty. Let (i_r, j_r) be the configuration at the beginning of round r . We say that the *Duplicator* has followed the *copy-cat* strategy in round r if $seg(i_{r+1}) - seg(i_r) = seg(j_{r+1}) - seg(j_r)$.

Proposition 9. *For a previously described n -round MTL[U, S] EF game on $M_{n,c,g}, N_{n,c,g}$ starting from $(1, 1)$, the *Duplicator* always has a winning strategy such that for each round $0 \leq r \leq n$, the i_r^{th} -event in $M_{n,c,g}$ and the j_r^{th} -event in $N_{n,c,g}$ satisfy the same set of propositions and*

⁷ We assume $n \cdot c' \geq \frac{1}{g} - 1$. If this is not the case, simply put more segments.




Fig. 1. A single segment in $N_{n,c,g}$




Fig. 2. Near the middle segments in $M_{n,c,g}$ and $N_{n,c,g}$

- if $seg(i_r) \neq seg(j_r)$, then

$$(n - r + 1) \cdot c' - 1 < seg(i_r), seg(j_r) < (n + r - 1) \cdot c' + 3.$$

Proof. We prove the proposition by induction on r . The basic idea is to make the resulting configurations identical whenever possible. Otherwise we try to adopt the copy-cat strategy. If that is also not possible, we choose another event that satisfies the same propositions. In the following, we refer to the timed word that the *Spoiler* first chooses as $\rho^s = (\sigma^s, \tau^s)$ ($\rho^d = (\sigma^d, \tau^d)$ for that of the *Duplicator*).

- *Base step.* The proposition holds trivially for $(i_0, j_0) = (1, 1)$.
- *Induction step.* Suppose the claim holds for $r < n$. We prove it also holds for $r + 1$.
 - $(i_r, j_r) = (1, 1)$:

Since all segments happen at time $> c \cdot g$, the *Duplicator* can always make (i_{r+1}, j_{r+1}) an identical configuration.
 - $(i_r, j_r) \neq (1, 1)$ is identical:

We may assume $r > 0$ and thus $n > 0$. Observe that any two p_1 -events that are $m \cdot (\frac{1}{g} - 1)$ segments away are separated by $m \cdot (1 - 2g)$, which is an integer multiple of g . More specifically, those p_1 -events whose distances to t_m are multiples of g will also have distances to t_l which are multiples of g . Consider the following cases:

 - * (i_r, j_r) both correspond to empty events and the *Spoiler* chooses event x or event y . However, the least distance between such a pair (i_r, j_r) and either event x or y is larger than $g + n \cdot c' \cdot (g - \frac{g}{\frac{1}{g} - 1}) + 2g - \frac{g}{\frac{1}{g} - 1} > c \cdot g$, which implies that the *Duplicator* can always make (i_{r+1}, j_{r+1}) identical.

- * (i_r, j_r) both correspond to p_1 -events and the *Spoiler* chooses (say) $i'_r = x$. The *Duplicator* may either choose $j'_r = y$ (then the *Duplicator* can surely make (i_{r+1}, j_{r+1}) identical later) or if that is not possible, choose event $j'_r = y'$. In the latter case, if the *Spoiler* plays **F**-part, it is obvious that the resulting configuration (i_{r+1}, j_{r+1}) would satisfy the claim. If she plays **U**-part, the *Duplicator* may either make (i_{r+1}, j_{r+1}) identical or $\text{seg}(j_{r+1}) - \text{seg}(i_{r+1}) = -1$. In this latter case it is clear that the claim still holds ($\text{seg}(i_{r+1}) = n \cdot c' + 2$ or $\text{seg}(i_{r+1}) = n \cdot c' + 4$). Notice that the assumption $g < \frac{1}{5}$ is crucial for the **U**-part (consider, e.g., $g = \frac{1}{5}$ and $\tau_{i_r}^s = \tau_{j_r}^d = t_l + g$).
- * (i_r, j_r) corresponds to p_2 -events except x and y , and the *Spoiler* chooses, say, event $i'_r = x$. The reasoning is exactly similar to the case above.
- * (i_r, j_r) corresponds to events x and y . If the *Spoiler* chooses some event z and forces the *Duplicator* not to choose the corresponding event but another one in a neighbouring segment, then that event z must be less than $(c + 1) \cdot g$ away from t_l . If it happens before t_l , then t_m would have distance $< (c - 1) \cdot g$ to it. If it happens after t_l , then t_m would be $< (c + 3) \cdot g$ away from it. Assume that z happens before t_l . If z is a p_1 -event, we divide $(c - 1) \cdot g$ by $(g - \frac{g}{\frac{1}{g}-1})$ to obtain $\frac{1-g}{1-2g} \cdot (c - 1) > |\text{seg}(z) - \text{seg}(i_r)|$ where $\text{seg}(i_r) = n \cdot c' + 1$. Observe that the p_1 -event z' that the *Duplicator* chooses as the response will be at most one more segment away. Then the claim holds regardless of the *Spoiler* plays **F**-part or **U**-part (may cause a drift of two more segments) later. If z is a p_2 -event, observe that its corresponding p_1 -event in the same segment must be at less than $2g + \frac{g}{\frac{1}{g}-1} < 3 \cdot (g - \frac{g}{\frac{1}{g}-1})$ away from z . Add this to $(c - 1) \cdot g$ and divide the result by $(g - \frac{g}{\frac{1}{g}-1})$ gives $\frac{1-g}{1-2g} \cdot (c - 1) + 3 < \frac{1-g}{1-2g} \cdot (c + 2)$. Again, the p_2 -event z' that the *Duplicator* chooses will at most one more segment away. The case for z happens after t_l is similar.

- (i_r, j_r) is not identical:

We claim that no matter how the *Spoiler* plays, the *Duplicator* can always either make (i_{r+1}, j_{r+1}) identical or, ensure that (i_{r+1}, j_{r+1}) has not moved towards the nearest end by $\geq c'$ segments. In the latter case the claim holds by IH. Note that we *cannot* always follow the copy-cat strategy. In the following, we would assume that the *Spoiler* always chooses some event that is more than two events away from the current event, e.g., $j'_r > j_r + 2$. If $j'_r \leq j_r + 2$, it is easy to see that the *Duplicator* can simply choose $i'_r = i_r + (j'_r - j_r)$.

Assume that (i_r, j_r) corresponds to a pair of p_1 -events and wlog. assume that the *Spoiler* chooses a position j'_r such that $j'_r > j_r$. If the *Duplicator* can choose i'_r such that $i'_r = j'_r$ and $i'_r \neq i_r + 1$, the *Duplicator* chooses $i'_r = j'_r$. Then, if the *Spoiler* plays **F**-part, it is immediate that $i_{r+1} = j_{r+1}$. If the *Spoiler* plays **U**-part, then the *Duplicator* makes $i_{r+1} = j_{r+1}$

whenever possible. Otherwise, for example, if $i_r < j_r$ and the *Spoiler* chooses some p_1 -event in $(\tau_{i_r}^d, \tau_{j_r}^d)$ as i_{r+1} , then the *Duplicator* chooses $j_{r+1} = j_r + 2$. Observe that i_{r+1} has moved towards j_r (and away from the nearest end). The claim holds by IH. If the *Duplicator* cannot choose i'_r such that $i'_r = j'_r$, or she can but $i'_r = i_r + 1$ prevents her from doing so, consider the following cases:

- * the *Duplicator* can choose i'_r such that $i'_r = i_r + (j'_r - j_r)$: If the *Duplicator* cannot choose $i'_r = j'_r$, then the *Duplicator* chooses $i'_r = i_r + (j'_r - j_r)$. As before, we know that $\tau_{j'_r}^s < \tau_{j_r}^s + (c+1) \cdot g$. It is easy to see that $\text{seg}(i_{r+1}) - \text{seg}(i_r) < c'$ and $\text{seg}(j_{r+1}) - \text{seg}(j_r) < c'$, and hence the claim holds by IH. If the *Duplicator* can indeed choose $i'_r = j'_r$ but $i'_r = i_r + 1$, then the *Duplicator* chooses $i'_r = i_r + 3$. It is easy to see that the claim holds by IH.
- * the *Duplicator* cannot choose i'_r such that $i'_r = i_r + (j'_r - j_r)$: This can only happen when j'_r corresponds to a p_2 -event. Observe that all p_1 -events in neighbouring segments are separated by $g - \frac{g}{\frac{1}{g}-1}$. The implication is that there exists t such that $t - \tau_{j_r}^s = m \cdot g = m' \cdot \frac{g}{\frac{1}{g}-1}$ for some $m, m' \in \mathbb{N}_{>0}$, and there exists $|k_1|, |k_2| < 1, k_1, k_2 \neq 0$ such that $t - \tau_{j_r}^s$ lies between
 - $\tau_{j'_r}^s - \tau_{j_r}^s = m_1 \cdot \frac{g}{\frac{1}{g}-1} + k_1 \cdot \epsilon, m_1 \in \mathbb{N}_{>0}$
 - $\tau_{i_r+(j'_r-j_r)}^d - \tau_{i_r}^d = m_2 \cdot \frac{g}{\frac{1}{g}-1} + k_2 \cdot \epsilon, m_2 \in \mathbb{N}_{>0}$.

It is obvious that $m_1 = m_2$. If $k_1 \cdot k_2 > 0$, since there is no integer multiple of $\frac{g}{\frac{1}{g}-1}$ that lies between, e.g., $m' \cdot \frac{g}{\frac{1}{g}-1}$ and $(m'+1) \cdot \frac{g}{\frac{1}{g}-1}$, this is a contradiction. If $k_1 \cdot k_2 < 0$, we must have $m' = m_1 = m_2$. This only happens when $i_r + (j'_r - j_r)$ in ρ^d corresponds to event x . In this case, the *Duplicator* chooses the corresponding event in a neighbouring segment. For example, if (i_r, j_r) corresponds to a pair of p_1 -events, $\text{seg}(i_r) = n \cdot c' + 1, \text{seg}(j_r) = n \cdot c', I = (2g, 3g)$ and $j'_r = y'$, then the *Duplicator* chooses $i'_r = x'$. Now if the *Spoiler* plays **F**-part, since we know that $\tau_{j'_r}^s < \tau_{j_r}^s + (c+1) \cdot g$, the claim holds. If the *Spoiler* plays **U**-part, e.g., in the fore-mentioned example, the *Spoiler* chooses $i_{r+1} = x$, then the *Duplicator* chooses $j_{r+1} = y''$. We see that the claim also holds.

Now assume that (i_r, j_r) corresponds to a pair of p_2 -events and wlog. assume that the *Spoiler* chooses a position j'_r such that $j'_r < j_r$. Most cases can be argued in very similar ways. We consider the situation when the *Duplicator* cannot choose i'_r such that $i'_r = i_r + (j'_r - j_r)$. If j'_r corresponds to a p_1 -event then the argument is exactly similar to above. Otherwise if j'_r corresponds to a p_2 -event, observe the fact that all p_2 -events in neighbouring segments, except x , are separated by $g - \frac{g}{\frac{1}{g}-1} + \frac{1}{3 \cdot n \cdot c' + 3} \cdot \epsilon$. By a similar argument, if $k_1 \cdot k_2 < 0$, the *Duplicator* chooses the corresponding event in a neighbouring segment. It can be argued in the same way that the claims holds regardless of the *Spoiler* plays **F**-part or **U**-part later.

□

Proposition 10. For a given MTL[U, S] formula φ of granularity g , modal depth n with largest constant $c \cdot g$, the constructed timed words $M_{n,c,g}, N_{n,c,g}$ satisfy

$$M_{n,c,g} \models \varphi \leftrightarrow N_{n,c,g} \models \varphi,$$

while

$$M_{n,c,g} \models \mathbf{F}(p_1 \wedge \mathfrak{E}_{(1,2)}^1 p_1) \text{ and } N_{n,c,g} \not\models \mathbf{F}(p_1 \wedge \mathfrak{E}_{(1,2)}^1 p_1).$$

Proof. Observe that for each p_1 -event in $N_{n,c,g}$, the first point in $(1, 2)$ may only be an empty event or a p_2 -event that is $\frac{1}{g} - 1$ segments later. In $M_{n,c,g}$, there will be a p_1 -event in $(1, 2)$ from some p_1 -event. The rest follows from Proposition 9 and Theorem 6. □

B Proof of Lemma 1

Proof. Since the satisfaction relations are defined inductively, we can work directly on the relevant subformula. We would like to prove that for a finite timed word ρ and a position i in ρ ,

$$\rho, i \models_f^+ \phi \leftrightarrow \rho, i \models_f^+ \phi' \text{ and } \rho, i \models_f^- \phi \leftrightarrow \rho, i \models_f^- \phi'$$

where $\phi \leftrightarrow \phi'$ matches one of the rules in Section 4. It is easy to see that the usual rules used to push negations or switch between modalities and their duals, e.g., $\mathbf{G}_I \varphi \leftrightarrow \neg \mathbf{F}_I \neg \varphi$, do not affect the informativeness of prefixes. For a group of similar rules we will only prove a representative one, as the proof for others follow similarly. In the following we let the LHS be ϕ and RHS be ϕ' .

$$- \varphi_1 \mathbf{U}_{(a,\infty)} \varphi_2 \leftrightarrow \varphi_1 \mathbf{U} \varphi_2 \wedge (\mathbf{F}_{(0,a]} \mathbf{true} \rightarrow \mathbf{G}_{(0,a]}(\varphi_1 \wedge \varphi_1 \mathbf{U} \varphi_2)):$$

$$\bullet \rho, i \models_f^+ \phi \leftrightarrow \rho, i \models_f^+ \phi':$$

Assume $\rho, i \models_f^+ \phi$. By definition we have $\rho, i \models_f^+ \varphi_1 \mathbf{U} \varphi_2$. If there is no event in $(\tau_i, \tau_i + a]$, since there must be an event in $(\tau_i + a, \tau_{|\rho|}]$, $\rho \models_f^- \mathbf{F}_{(0,a]} \mathbf{true}$ cannot hold and we are done. If there are events in $(\tau_i, \tau_i + a]$, then for all j such that $\tau_j - \tau_i \in (0, a]$ we have $\rho, j \not\models_f^- \neg \varphi_1$. Also for all such j we have $\rho, j \not\models_f^- \neg \varphi_1 \mathbf{U} \varphi_2$ since it is obvious that $\rho, j \models_f^+ \varphi_1 \mathbf{U} \varphi_2$. For the other direction, if $\rho, i \models_f^+ \varphi_1 \mathbf{U} \varphi_2$ because of a witness in $(\tau_i + a, \tau_{|\rho|})$ then we are done. If this is not the case we must have $\rho, i \models_f^- \mathbf{F}_{(0,a]} \mathbf{true}$ (since there is an event in $(\tau_i, \tau_i + a]$). We then have $\rho, i \not\models_f^- \mathbf{F}_{(0,a]}(\neg \varphi_1 \vee \neg(\varphi_1 \mathbf{U} \varphi_2))$ (this implies $\tau_{|\rho|} \geq a$). Now for all j such that $\tau_j - \tau_i \in (0, a]$ we have $\rho, j \models_f^+ \varphi_1$ and $\rho, j \models_f^+ \varphi_1 \mathbf{U} \varphi_2$, which imply $\rho, i \models_f^+ \phi$.

$$\bullet \rho, i \models_f^- \phi \leftrightarrow \rho, i \models_f^- \phi':$$

Assume $\rho, i \models_f^- \phi$. This holds if there is a witness in (a, ∞) or $\rho, i \models_f^- \mathbf{G} \varphi_1$. In both cases we have $\rho, i \models_f^- \varphi_1 \mathbf{U} \varphi_2$. If there is no event in $(\tau_i, \tau_i + a]$ then

we are done. If there is a witness, then for all such j that $\tau_j - \tau_i \in (0, a]$ we have $\rho, j \models_{\mathcal{F}} \varphi_1$ and $\rho, j \models_{\mathcal{F}} \varphi_1 \mathbf{U} \varphi_2$. If there is no witness then for all such j we again have $\rho, j \models_{\mathcal{F}} \varphi_1$ and $\rho, j \models_{\mathcal{F}} \varphi_1 \mathbf{U} \varphi_2$. For the other direction, if there is no event in $(\tau_i, \tau_i + a]$ we are done. If there are event in $(\tau_i, \tau_i + a]$, all j such that $\tau_j - \tau_i \in (0, a]$ will satisfy $\rho, j \models_{\mathcal{F}} \varphi_1$ and $\rho, j \models_{\mathcal{F}} \varphi_1 \mathbf{U} \varphi_2$. This clearly gives $\rho, i \models_{\mathcal{F}} \phi$.

$$- \neg(\varphi_1 \mathbf{U} \varphi_2) \leftrightarrow \mathbf{G}\neg\varphi_2 \vee (\neg\varphi_2 \mathbf{U} (\neg\varphi_2 \wedge \neg\varphi_1)):$$

$$\bullet \rho, i \models_{\mathcal{F}}^+ \phi \leftrightarrow \rho, i \models_{\mathcal{F}}^+ \phi':$$

Assume $\rho, i \models_{\mathcal{F}}^+ \phi \leftrightarrow \rho, i \not\models_{\mathcal{F}}^+ \varphi_1 \mathbf{U} \varphi_2$. This implies that φ_1 fails to hold before φ_2 holds, and we have $\rho, i \models_{\mathcal{F}}^+ \neg\varphi_2 \mathbf{U} (\neg\varphi_2 \wedge \neg\varphi_1)$. For the other direction note that $\rho, i \not\models_{\mathcal{F}}^+ \mathbf{G}\neg\varphi_2$, the second disjunct must be satisfied, and it is easy to see that $\rho, i \models_{\mathcal{F}}^+ \phi$.

$$\bullet \rho, i \models_{\mathcal{F}}^- \phi \leftrightarrow \rho, i \models_{\mathcal{F}}^- \phi':$$

Assume $\rho, i \models_{\mathcal{F}}^- \neg(\varphi_1 \mathbf{U} \varphi_2) \leftrightarrow \rho, i \not\models_{\mathcal{F}}^+ \varphi_1 \mathbf{U} \varphi_2$. This implies either $\rho, j \not\models_{\mathcal{F}}^+ \varphi_2 \leftrightarrow \rho, j \models_{\mathcal{F}}^- \neg\varphi_2$ for all $j > i$ in ρ (this gives $\rho, i \models_{\mathcal{F}}^- \mathbf{G}\neg\varphi_2$) or φ_1 fails to hold before φ_2 holds— $\rho, i \models_{\mathcal{F}}^- \neg\varphi_2 \mathbf{U} (\neg\varphi_2 \wedge \neg\varphi_1)$. For the other direction, if $\rho, i \models_{\mathcal{F}}^- \mathbf{G}\neg\varphi_2 \leftrightarrow \rho, i \not\models_{\mathcal{F}}^+ \mathbf{F}\varphi_2$ then $\rho, i \models_{\mathcal{F}}^+ \varphi_1 \mathbf{U} \varphi_2$ cannot hold. If $\rho, i \models_{\mathcal{F}}^- \neg\varphi_2 \mathbf{U} (\neg\varphi_2 \wedge \neg\varphi_1)$ then either $\rho, i \models_{\mathcal{F}}^- \mathbf{G}\neg\varphi_2$ or there is a witness, and it is easy to see that $\rho, i \models_{\mathcal{F}}^+ \varphi_1 \mathbf{U} \varphi_2$ cannot hold.

$$- \mathfrak{E}_{(a,\infty)}^c \varphi \leftrightarrow \mathfrak{E}_{(a,a+\max(c,a-c))}^c \varphi \vee (\neg \mathbf{F}_{(c,a)} \mathbf{true} \wedge \mathbf{F}_{[0,c]}^w (\mathbf{false} \mathbf{U}_{(c,\infty)} \varphi)) \text{ (let } c \geq a - c \text{):}$$

$$\bullet \rho, i \models_{\mathcal{F}}^+ \phi \leftrightarrow \rho, i \models_{\mathcal{F}}^+ \phi':$$

Assume $\rho, i \models_{\mathcal{F}}^+ \phi$. If there is an event in $(\tau_i + c, \tau_i + a]$ then since $c \geq a - c$, the witness must be in $(\tau_i + a, \tau_i + a + c]$, and clearly $\rho, i \models_{\mathcal{F}}^+ \mathfrak{E}_{(a,a+c]}^c \varphi$. If there is no event in $(\tau_i + c, \tau_i + a]$ then $\rho, i \models_{\mathcal{F}}^- \mathbf{F}_{(c,a)} \mathbf{true}$ cannot hold. In this case, if the witness is in $(\tau_i + a, \tau_i + a + c]$ then $\rho, i \models_{\mathcal{F}}^+ \mathfrak{E}_{(a,a+c]}^c \varphi$. If the witness is in $(\tau_i + a + c, \tau_{|\rho|}]$ then $\rho, i \models_{\mathcal{F}}^+ \mathbf{F}_{[0,c]}^w (\mathbf{false} \mathbf{U}_{(c,\infty)} \varphi)$ as $a \geq c$. For the other direction, if $\rho, i \models_{\mathcal{F}}^+ \mathfrak{E}_{(a,a+c]}^c \varphi$ then clearly $\rho, i \models_{\mathcal{F}}^+ \phi$. If $\rho, i \not\models_{\mathcal{F}}^+ \mathbf{F}_{(c,a)} \mathbf{true}$, note that the witness for $\rho, i \models_{\mathcal{F}}^+ \mathbf{F}_{[0,c]}^w (\mathbf{false} \mathbf{U}_{(c,\infty)} \varphi)$ must be the last event in $[\tau_i, \tau_i + c]$, and it is clear that $\rho, i \models_{\mathcal{F}}^+ \phi$.

$$\bullet \rho, i \models_{\mathcal{F}}^- \phi \leftrightarrow \rho, i \models_{\mathcal{F}}^- \phi':$$

Assume $\rho, i \models_{\mathcal{F}}^- \phi$. If $\tau_{|\rho|} \leq \tau_i + a$ then $\rho, i \models_{\mathcal{F}}^- \mathfrak{E}_{(a,a+c]}^c \varphi$. If $\tau_{|\rho|} > \tau_i + a$ and the witness is in $(\tau_i + a, \tau_i + a + c]$ then again $\rho, i \models_{\mathcal{F}}^- \mathfrak{E}_{(a,a+c]}^c \varphi$. Otherwise $\tau_{|\rho|} > \tau_i + a + c$ and there is no event in $(\tau_i + a, \tau_i + a + c]$. This implies that there is no event in $(\tau_i + c, \tau_i + a]$ and hence $\rho, i \not\models_{\mathcal{F}}^+ \mathbf{F}_{(c,a)} \mathbf{true}$. Since $\tau_{|\rho|} > \tau_i + a + c$, $\rho, i \models_{\mathcal{F}}^- \mathbf{F}_{[0,c]}^w (\mathbf{false} \mathbf{U}_{(c,\infty)} \varphi)$ clearly holds. For the other direction, if $\rho, i \models_{\mathcal{F}}^- \mathfrak{E}_{(a,a+c]}^c \varphi$ then clearly $\rho, i \models_{\mathcal{F}}^- \phi$.

If $\rho, i \not\models_f^+ \mathbf{F}_{(c,a)} \mathbf{true}$, then if $\tau_{|\rho|} \leq c$ we are done. Otherwise if $\tau_{|\rho|} > a$, since $\rho, i \models_f^- \mathbf{F}_{[0,c]}^w(\mathbf{false} \mathbf{U}_{(c,\infty)} \varphi)$, the first event in $(\tau_i + a, \tau_{|\rho|}]$ must be a witness for $\rho, i \models_f^- \phi$.

– $\mathfrak{E}_{(a,b)}^c(\varphi_1 \mathbf{U} \varphi_2) \leftrightarrow (\mathfrak{E}_{(a,b)}^c \mathbf{true})$

$\wedge \left(\mathfrak{E}_{(a,b)}^a(\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2) \vee (\mathfrak{E}_{(a,b)}^a(\mathbf{G}_{(0,2b)} \varphi_1) \wedge \varphi_{ugb}) \right)$:

• $\rho, i \models_f^+ \phi \leftrightarrow \rho, i \models_f^+ \phi'$:

Assume $\rho, i \models_f^+ \phi$. It is clear that $\rho, i \models_f^+ \mathfrak{E}_{(a,b)}^c \mathbf{true}$. With an argument similar to the proof of Proposition 5, it is easy to see that $\rho, i \models_f^+ \mathfrak{E}_{(a,b)}^a(\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2)$ or $\rho, i \models_f^+ \mathfrak{E}_{(a,b)}^a(\mathbf{G}_{(0,2b)} \varphi_1) \wedge \varphi_{ugb}$. For the other direction, if $\rho, i \models_f^+ \mathfrak{E}_{(a,b)}^a(\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2)$ then we are done. Otherwise, the argument is similar to the proof of Proposition 5.

• $\rho, i \models_f^- \phi \leftrightarrow \rho, i \models_f^- \phi'$:

Assume $\rho, i \models_f^- \phi$. If $\tau_{|\rho|} \leq a$ then we are done. Otherwise if $\tau_{|\rho|} > a$, we must have $\rho, i \models_f^- \mathfrak{E}_{(a,b)}^c \mathbf{true}$. The rest is similar to the proof of Proposition 5. The other direction is also similar to the proof of Proposition 5.

– $((\varphi_1 \mathbf{U} \varphi_2) \vee \chi) \mathbf{U}_{(a,b)} \theta \leftrightarrow ((\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2) \vee \chi) \mathbf{U}_{(a,b)} \theta$

$\vee \left(\left((\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2) \vee \chi \right) \mathbf{U}_{(0,b)} (\mathbf{G}_{(0,2b)} \varphi_1) \wedge \mathbf{F}_{(a,b)} \theta \wedge \varphi_{ugb} \right)$:

• $\rho, i \models_f^+ \phi \leftrightarrow \rho, i \models_f^+ \phi'$:

Assume $\rho, i \models_f^+ \phi$. It is obvious that $\rho, i \models_f^+ \mathbf{F}_{(a,b)} \theta$ holds. If the first disjunct of ϕ' does not hold, then $\rho, i \models_f^+ ((\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2) \vee \chi) \mathbf{U}_{(0,b)} (\mathbf{G}_{(0,2b)} \varphi_1)$ must hold. The last conjunct holds by an argument similar to the proof of Proposition 5. For the other direction, if the first disjunct of ϕ' holds then we are done. If it does not hold, then there must be a witness (at which φ_2 holds) in $[\tau_i + 2b, \tau_{|\rho|}]$, and it is easy to see that $\rho, i \models_f^+ \phi$.

• $\rho, i \models_f^- \phi \leftrightarrow \rho, i \models_f^- \phi'$:

Assume $\rho, i \models_f^- \phi$. If the first disjunct of ϕ' does not hold then there must be events in $[\tau_i + 2b, \tau_{|\rho|}]$. It follows that $\rho, i \models_f^- ((\varphi_1 \mathbf{U}_{(0,2b)} \varphi_2) \vee \chi) \mathbf{U}_{(0,b)} (\mathbf{G}_{(0,2b)} \varphi_1)$ and $\rho, i \models_f^- \mathbf{F}_{(a,b)} \theta$ must hold. The rest is similar to the proof to Proposition 5. For the other direction, if the first disjunct of ϕ' holds then we are done. Otherwise if $\tau_{|\rho|} < b$, it is easy to see that $\rho, i \models_f^- \phi$. If this is not the case then the proof closely follows Proposition 5. \square

C Proof of Remark 1

Consider Safety-MTL with past: when written in *negation normal form*, all \mathbf{U}_I are bounded while there are no restrictions on \mathbf{S}_I .

Proposition 11 (Safety). $\eta, i \not\models \varphi$ implies that there is a prefix ρ of η s.t. for any ζ , $\rho\zeta, i \not\models \varphi$.

Proof. Simple induction on the structure of φ . \square

Proposition 12 (Non-pathological safety). For a prefix ρ , if $\rho\zeta, i \not\models \varphi$ for all ζ , then for any ζ there exists a prefix ρ' of $\rho\zeta$ s.t. $\rho', i \models_f^+ \neg\varphi$.

Proof. We prove this by structural induction. The base step is trivial. For the induction step:

- $\varphi_1 \vee \varphi_2$:
For arbitrary ζ we have $\rho\zeta, i \not\models \varphi_1$ and $\rho\zeta, i \not\models \varphi_2$. Hence for any ζ there is a ρ' such that $\rho', i \models_f^+ \neg\varphi_1$ and $\rho', i \models_f^+ \neg\varphi_2$ (by IH and Proposition 4). We then have $\rho', i \models_f^+ \neg\varphi_1 \wedge \rho', i \models_f^+ \neg\varphi_2 \leftrightarrow \rho', i \models_f^+ (\neg\varphi_1 \wedge \neg\varphi_2) \leftrightarrow \rho', i \models_f^+ \neg(\varphi_1 \vee \varphi_2)$.
- $\varphi_1 \wedge \varphi_2$:
For arbitrary ζ we have $\rho\zeta, i \not\models \varphi_1$ and $\rho\zeta, i \not\models \varphi_2$. Hence for any ζ there is a ρ' such that $\rho', i \models_f^+ \neg\varphi_1$ or $\rho', i \models_f^+ \neg\varphi_2$ (by IH). We then have $\rho' \not\models_f^+ \varphi_1 \vee \rho' \not\models_f^+ \varphi_2 \leftrightarrow \rho' \not\models_f^+ \varphi_1 \wedge \varphi_2 \leftrightarrow \rho' \models_f^+ \neg(\varphi_1 \wedge \varphi_2)$.
- $\varphi_1 \mathbf{U}_I \varphi_2$:
 $\rho\zeta, i \not\models \varphi_1 \mathbf{U}_I \varphi_2$ for any ζ iff for any ζ , φ_1 fails to hold before φ_2 holds or there is no φ_2 (later than τ_i) in I . By non-Zenoness and the fact that I is bounded, in the first case we have, for a given ζ , a prefix ρ' of $\rho\zeta$ s.t. $\exists j \in (i, |\rho'|] \left(\tau_j - \tau_i \leq I \wedge \rho', j \models_f^+ \neg\varphi_1 \wedge \forall l \in (i, j] (\rho', l \models_f^+ \neg\varphi_2) \right)$ by IH. In the second case we have ρ' s.t. $\forall j \in (i, |\rho'|] (\tau_j - \tau_i \in I \rightarrow \rho', j \models_f^+ \neg\varphi_2)$. In either case it is easy to see that $\rho', i \models_f^+ \neg(\varphi_1 \mathbf{U}_I \varphi_2)$.
- $\varphi_1 \tilde{\mathbf{U}}_I \varphi_2$: $\rho\zeta, i \not\models \varphi_1 \tilde{\mathbf{U}}_I \varphi_2$ for arbitrary ζ iff for any ζ , we have $\neg\varphi_1 \mathbf{U}_I \neg\varphi_2$ holds at i . We have, for a given ζ , a prefix ρ' of $\rho\zeta$ such that $\exists j \in (i, |\rho'|] \left(\tau_j - \tau_i \in I \wedge \rho', j \models_f^+ \neg\varphi_2 \wedge \forall l \in (i, j) (\rho', l \models_f^+ \neg\varphi_1) \right)$ by IH. Now, of course, $\rho', i \models_f^+ \neg\varphi_1 \mathbf{U}_I \neg\varphi_2 \leftrightarrow \rho', i \models_f^+ \neg(\varphi_1 \tilde{\mathbf{U}}_I \varphi_2)$.

The cases for $\varphi_1 \mathbf{S}_I \varphi_2$ and $\varphi_1 \tilde{\mathbf{S}}_I \varphi_2$ are exactly similar. \square