

On the Expressiveness of Metric Temporal Logic over Bounded Timed Words

Hsi-Ming Ho and Joël Ouaknine

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

Abstract. It is known that Metric Temporal Logic (MTL) is strictly less expressive than the Monadic First-Order Logic of Order and Metric ($\text{FO}[\langle, +1]$) in the pointwise semantics over bounded time domains (i.e., timed words of bounded duration) [2]. In this paper, we present an extension of MTL, relying on two new temporal operators ‘Earlier’ (\mathfrak{E}) and ‘Newer’ (\mathfrak{N}), which has the same expressive power as $\text{FO}[\langle, +1]$ in both the pointwise and continuous semantics over bounded time domains.

1 Introduction

One of the most prominent specification formalisms used in verification is *Linear Temporal Logic* (LTL), which is typically interpreted over the non-negative integers or reals. A celebrated result of Kamp [7] states that, in either case, LTL has precisely the same expressive power as the *Monadic First-Order Logic of Order* ($\text{FO}[\langle]$). These logics, however, are inadequate to express specifications for systems whose correct behaviour depends on quantitative timing requirements. Over the last three decades, much work has therefore gone into lifting classical verification formalisms and results to the real-time setting. *Metric Temporal Logic* (MTL), which extends LTL by constraining the temporal operators by time intervals, was introduced by Koymans [8] in 1990 and has emerged as a central real-time specification formalism.

MTL enjoys two main semantics, depending intuitively on whether atomic formulas are interpreted as *state predicates* or as (instantaneous) *events*. In the former, the system is assumed to be under observation at every instant in time, leading to a ‘continuous’ semantics based on *flows* or *signals*, whereas in the latter observations of the system are taken to be (finite or infinite) sequences of timestamped snapshots, leading to a ‘pointwise’ semantics based on *timed words*—this is the leading interpretation, for example, for systems modelled as timed automata. In both cases, the time domain is usually taken to be the non-negative real numbers. Both semantics have been extensively studied; see, e.g., [10] for a partial historical account.

Alongside these developments, researchers proposed the *Monadic First-Order Logic of Order and Metric* ($\text{FO}[\langle, +1]$) as a natural quantitative extension of $\text{FO}[\langle]$. Unfortunately, Hirshfeld and Rabinovich [4] showed that no ‘finitary’ extension of MTL—and *a fortiori* MTL itself—could have the same expressive

power as $\text{FO}[<, +1]$ over the non-negative reals, shattering the hope of lifting Kamp’s theorem to the real-time world.¹ The situation was however partly improved recently by showing that, over *bounded* time domains, MTL has precisely the same expressive power as $\text{FO}[<, +1]$ in the continuous semantics [9, 11].² Nonetheless, and rather surprisingly, MTL remains strictly less expressive than $\text{FO}[<, +1]$ over bounded time domains in the pointwise semantics, i.e., over timed words of bounded duration, as can immediately be seen from D’Souza and Prabhakar’s construction [2].

The main result of this paper is to show that MTL, equipped with both the forwards and backwards temporal modalities ‘Until’ (\mathcal{U}) and ‘Since’ (\mathcal{S}), together with two new modalities ‘Earlier’ (\mathcal{E}) and ‘Newer’ (\mathcal{N}), has precisely the same expressive power as $\text{FO}[<, +1]$ over bounded time domains in the pointwise semantics (and also, trivially, in the continuous semantics). This augmented version of Metric Temporal Logic, written $\text{MTL}[\mathcal{U}, \mathcal{S}, \mathcal{E}, \mathcal{N}]$, therefore yields a definitive real-time analogue of Kamp’s theorem over bounded domains.

It is worth noting that $\text{MTL}[\mathcal{U}, \mathcal{S}, \mathcal{E}, \mathcal{N}]$ satisfiability and model checking (against timed automata) are decidable over bounded time domains, thanks to the decidability of $\text{FO}[<, +1]$ over such domains as established in [9, 11]. Unfortunately, $\text{FO}[<, +1]$ has non-elementary complexity, whereas the time-bounded satisfiability and model-checking problems for MTL are EXPSPACE-complete [9, 11]. However, it can easily be seen by inspecting the relevant constructions that the complexity bounds for MTL carry over to our new logic $\text{MTL}[\mathcal{U}, \mathcal{S}, \mathcal{E}, \mathcal{N}]$.

2 Preliminaries

2.1 Timed Words

A *time sequence* $\tau = \tau_1\tau_2\dots$ is a non-empty finite sequence over non-negative reals that satisfies the requirements below (we denote the length of τ by $|\tau|$):

- *Initialisation*: $\tau_1 = 0$
- *Strict monotonicity*: For all $i, 1 \leq i < |\tau|$, we have $\tau_i < \tau_{i+1}$.

A *timed word* over finite alphabet Σ is a pair $\rho = (\sigma, \tau)$, where $\sigma = \sigma_1\sigma_2\dots$ is a non-empty finite word over Σ and τ is a time sequence of the same length. We refer to each (σ_i, τ_i) as an *event*. In this sense, a timed word can be regarded as a sequence of events. We denote by $|\rho|$ the number of events in ρ . A *position* in ρ is a number i such that $1 \leq i \leq |\rho|$. The *duration* of ρ is defined as $\tau_{|\rho|}$.

¹ Hirshfeld and Rabinovich’s result was only stated and proved for the continuous semantics, but we believe that their approach would also carry through for the pointwise semantics. In any case, using different techniques D’Souza and Prabhakar [2] independently showed that MTL is strictly weaker than $\text{FO}[<, +1]$ in the pointwise semantics.

² See also [5, 6] which present further expressive completeness results for extensions of MTL over the reals in the continuous semantics.

Fix a positive integer N . We denote by \mathbb{T} the bounded interval $[0, N)$ or $\mathbb{R}_{\geq 0}$. A \mathbb{T} -timed word is a timed word with all its timestamps in \mathbb{T} . We write $\mathbb{T}\Sigma^*$ for the set of (finite) \mathbb{T} -timed words over Σ . For a set of monadic predicates \mathbf{P} we write $\Sigma_{\mathbf{P}} = 2^{\mathbf{P}}$.

Note that we are focussing on finite timed words. Our results carry over to the case of (Zeno) infinite timed words as well, with some modifications.

2.2 Metric Logics

We first define a metric predicate logic $\text{FO}[\langle, +1]$ which will serve as a ‘yardstick’ of expressiveness in this paper.

Definition 1. *Given a set of monadic predicates \mathbf{P} , the set of $\text{FO}[\langle, +1]$ formulas is generated by the grammar*

$$\vartheta ::= P(x) \mid x < x' \mid d(x, x') \approx c \mid \mathbf{true} \mid \vartheta_1 \wedge \vartheta_2 \mid \neg \vartheta \mid \exists x \vartheta,$$

where $P \in \mathbf{P}$, x, x' are variables, $\approx \in \{=, \neq, <, >, \leq, \geq\}$ and $c \in \mathbb{N}_{\geq 0}$.

With each \mathbb{T} -timed word $\rho = (\sigma, \tau)$ over $\Sigma_{\mathbf{P}}$ we associate a structure M_ρ . Its universe U_ρ is the finite subset $\{\tau_i \mid 1 \leq i \leq |\rho|\}$ of \mathbb{T} . The order relation $<$ and monadic predicates in \mathbf{P} are interpreted in the expected way. The binary *distance predicate* $d(x, x') \approx c$ holds iff $|x - x'| \approx c$. The satisfaction relation is defined inductively as usual. We write $M_\rho, t_1, \dots, t_n \models \vartheta(x_1, \dots, x_n)$ (or $\rho, t_1, \dots, t_n \models \vartheta(x_1, \dots, x_n)$) if $t_1, \dots, t_n \in U_\rho$ and $\vartheta(t_1, \dots, t_n)$ holds in M_ρ . We say that $\text{FO}[\langle, +1]$ formulas $\vartheta_1(x)$ and $\vartheta_2(x)$ are *equivalent* over \mathbb{T} -timed words if for all \mathbb{T} -timed words ρ and $t \in U_\rho$,

$$\rho, t \models \vartheta_1(x) \iff \rho, t \models \vartheta_2(x).$$

Formulas of metric temporal logics are built from monadic predicates using Boolean connectives and *modalities*. A k -ary modality is defined by an $\text{FO}[\langle, +1]$ formula $\varphi(x, X_1, \dots, X_k)$ with free second-order variables X_1, \dots, X_k and free first-order variable x . For example, the MTL modality $\mathcal{U}_{(0,5)}$ is defined by

$$\begin{aligned} \mathcal{U}_{(0,5)}(x, X_1, X_2) = \exists x' \left(x < x' \wedge d(x, x') < 5 \wedge X_2(x') \right. \\ \left. \wedge \forall x'' (x < x'' \wedge x'' < x' \implies X_1(x'')) \right). \end{aligned}$$

In MTL notation, this is written using infix notation $X_1 \mathcal{U}_{(0,5)} X_2$.

Definition 2. *Given a set of monadic predicates \mathbf{P} , the set of $\text{MTL}[\mathcal{U}, \mathcal{S}]$ formulas is generated by the grammar*

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

where $P \in \mathbf{P}$ and $I \subseteq (0, \infty)$ is an interval with endpoints in $\mathbb{N}_{\geq 0} \cup \{\infty\}$.

For each I , let $l(I) = \inf(I)$, $r(I) = \sup(I)$ and $|I| = r(I) - l(I)$. If I is not present as a subscript to a given modality then it is assumed to be $(0, \infty)$. We freely use the usual syntactic sugar, e.g., $\mathbf{false} \equiv \neg \mathbf{true}$, $\diamond_I \varphi \equiv \mathbf{true} \mathcal{U}_I \varphi$, $\square_I \varphi \equiv \neg \diamond_I \neg \varphi$ and $\bigcirc_I \varphi \equiv \mathbf{false} \mathcal{U}_I \varphi$. For the sake of completeness, we give a traditional inductive definition of the satisfaction relation of $\text{MTL}[\mathcal{U}, \mathcal{S}]$ below.

Definition 3. *The satisfaction relation $(\rho, i) \models \varphi$ for an $\text{MTL}[\mathcal{U}, \mathcal{S}]$ formula φ , a timed word $\rho = (\sigma, \tau)$ and a position i in ρ is defined as follows:*

- $(\rho, i) \models P$ iff $P(\tau_i)$ holds in M_ρ
- $(\rho, i) \models \mathbf{true}$
- $(\rho, i) \models \varphi_1 \wedge \varphi_2$ iff $(\rho, i) \models \varphi_1$ and $(\rho, i) \models \varphi_2$
- $(\rho, i) \models \neg \varphi$ iff $(\rho, i) \not\models \varphi$
- $(\rho, i) \models \varphi_1 \mathcal{U}_I \varphi_2$ iff there exists $j, i < j \leq |\rho|$ such that $(\rho, j) \models \varphi_2$, $\tau_j - \tau_i \in I$, and $(\rho, k) \models \varphi_1$ for all k with $i < k < j$
- $(\rho, i) \models \varphi_1 \mathcal{S}_I \varphi_2$ 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 $j < k < i$.

Note that we adopt strict versions of \mathcal{U}_I and \mathcal{S}_I . For convenience, we also define weak versions of these modalities, e.g., $\varphi_1 \mathcal{U}_I^w \varphi_2 \equiv \varphi_1 \wedge (\varphi_1 \mathcal{U}_I \varphi_2)$ if $0 \notin I$ and $\varphi_1 \mathcal{U}_I^w \varphi_2 \equiv \varphi_2 \vee (\varphi_1 \wedge (\varphi_1 \mathcal{U}_I \varphi_2))$ if $0 \in I$ (in this case we allow $0 \in I$ and assume $I = [0, \infty)$ if the subscript is absent). We write $\rho \models \varphi$ if $\rho, 1 \models \varphi$.

2.3 Relative Expressiveness

Let L, L' be two metric logics. We say that L' is **at least as expressive as** L ($L \subseteq L'$) over \mathbb{T} -timed words if, for any sentence ϑ of L , there exists a sentence φ of L' such that ϑ and φ are satisfied by precisely the same set of \mathbb{T} -timed words. They are said to be **equally expressive** ($L = L'$) over \mathbb{T} -timed words if each is at least as expressive as the other over \mathbb{T} -timed words. When L' is a temporal logic and L is a predicate logic, we say that L' is **expressively complete for** L ($L \stackrel{\leftrightarrow}{=} L'$) over \mathbb{T} -timed words if for any formula $\vartheta(x) \in L$, there is an equivalent formula $\varphi \in L'$ over \mathbb{T} -timed words.

3 Expressiveness

In this section, we present a sequence of successively more expressive extensions of MTL over bounded timed words, culminating in the next section with a logic that is expressive complete. Several of the modalities introduced below will feature in the proof of our main result. Along the way we highlight the key requirements for expressive completeness.

3.1 Definability of Time 0

Recall that MTL and $\text{FO}[<, +1]$ have the same expressiveness over continuous domains of the form $[0, N)$ [9], a result that fails over $[0, N)$ -timed words. To

account for this difference between the two semantics, observe that a distinctive feature of the continuous semantics is exploited in [9]: in any $[0, N)$ -flow, the formula $\diamond_{=(N-1)} \mathbf{true}$ holds in $[0, 1)$ and nowhere else. One can make use of conjunctions of similar formulas to tell which unit interval the current instant is in. Apparently, this trick does not work for MTL in the pointwise semantics. However, for $\text{MTL}[\mathcal{U}, \mathcal{S}]$, this ability can be achieved by using past modalities. We define

$$\varphi_{i,i+1} = \diamond_{[i,i+1)}(\neg \diamond \mathbf{true})$$

and let $\Phi_{unit} = \{\varphi_{i,i+1} \mid i \in \mathbb{N}_{\geq 0}\}$. It is clear that $\varphi_{i,i+1}$ holds only in $[i, i+1)$ and nowhere else. Indeed, our main result depends crucially on the use of these formulas. Denote by $\text{MTL}[\Phi_{unit}]$ the extension of MTL obtained by allowing these formulas as subformulas. This very restrictive use of past modalities strictly increases the expressiveness of MTL.

Proposition 1. $\text{MTL} \subsetneq \text{MTL}[\Phi_{unit}]$ over $[0, N)$ -timed words.

Proof. For a given $m \in \mathbb{N}_{\geq 0}$, we construct the following models:

$$\begin{aligned} \mathcal{A}_m &= (\emptyset, 0)(\emptyset, 1 - \frac{2.5}{2m+5})(\emptyset, 1 - \frac{1.5}{2m+5})(\emptyset, 1 - \frac{0.5}{2m+5}) \dots (\emptyset, 1 + \frac{m+2.5}{2m+5}), \\ \mathcal{B}_m &= (\emptyset, 0)(\emptyset, 1 - \frac{1.5}{2m+5})(\emptyset, 1 - \frac{0.5}{2m+5})(\emptyset, 1 + \frac{0.5}{2m+5}) \dots (\emptyset, 1 + \frac{m+3.5}{2m+5}). \end{aligned}$$




Fig. 1. Models \mathcal{A}_m and \mathcal{B}_m

The models are illustrated in Figure 1 where each hollow box represents an event. It can be shown that no MTL formula of modal depth $\leq m$ distinguishes \mathcal{A}_m and \mathcal{B}_m while there is a simple $\text{MTL}[\Phi_{unit}]$ formula does. We detail the proof in Appendix. \square

This conservative extension is not sufficient, however: the following proposition says that non-trivial nesting of future modalities and past modalities gives more expressiveness.

Proposition 2. $\text{MTL}[\Phi_{unit}] \subsetneq \text{MTL}[\mathcal{U}, \mathcal{S}]$ over $[0, N)$ -timed words.

Proof. For a given $m \in \mathbb{N}_{\geq 0}$, we construct

$$\mathcal{C}_m = (\emptyset, 0)(\emptyset, 0.5 + \frac{0.5}{4m+6})(\emptyset, 0.5 + \frac{1.5}{4m+6}) \dots (\emptyset, 1.5 - \frac{0.5}{4m+6}).$$




Fig. 2. Models \mathcal{C}_m and \mathcal{D}_m

\mathcal{D}_m is constructed as \mathcal{C}_m except that the event at time $0.5 + \frac{m+1.5}{4m+6}$ is missing.

The models are illustrated in Figure 2 where each hollow box represents an event. It can be shown that no MTL[Φ_{unit}] formula of modal depth $\leq m$ distinguishes \mathcal{C}_m and \mathcal{D}_m while there is a simple MTL[\mathcal{U}, \mathcal{S}] formula does. We give detailed proof in Appendix. \square

3.2 Counting Modalities

The modality $C_n(x, \varphi)$ asserts that φ holds at least at n points in the open interval $(x, x+1)$. The modalities C_n for $n \geq 2$ are called the *counting modalities*. It is well-known that these modalities are inexpressible in MTL[\mathcal{U}, \mathcal{S}] over $\mathbb{R}_{\geq 0}$ -flows [3]. For this reason, they (or variants thereof) are often used to separate the expressiveness of various metric logics (cf., e.g., [1, 12, 13]). For example, the FO[$<, +1$] formula

$$\vartheta_{pqr} = \exists x \left(P(x) \wedge \exists x' \left(x < x' \wedge d(x, x') > 1 \wedge d(x, x') < 2 \wedge Q(t_1) \right. \right. \\ \left. \left. \wedge \exists x'' \left(x' < x'' \wedge d(x, x'') > 1 \wedge d(x, x'') < 2 \wedge R(x'') \right) \right) \right)$$

has no equivalent in MTL[\mathcal{U}, \mathcal{S}] over $\mathbb{R}_{\geq 0}$ -timed words [12]. Indeed, it is shown very recently that in the continuous semantics, MTL[\mathcal{U}, \mathcal{S}] with counting modalities and their past versions (which we denote by MTL[$\mathcal{U}, \mathcal{S}, \{C_n, \overleftarrow{C}_n\}_{n=2}^{\infty}$]) is expressively complete for FO[$<, +1$] [5]. However, counting modalities are expressible in MTL[\mathcal{U}, \mathcal{S}] in the time-bounded setting. To see this, observe that the following formula is equivalent to ϑ_{pqr} over $[0, N)$ -timed words (we make use of formulas in Φ_{unit} defined in the last subsection):

$$\diamond \left(\bigvee_{i \in [0, N-1]} \left(P \wedge \varphi_{i, i+1} \wedge \left(\diamond_{>1} (Q \wedge \diamond (R \wedge \varphi_{i+1, i+2})) \right. \right. \right. \\ \left. \left. \vee \diamond_{<2} (R \wedge \varphi_{i+2, i+3} \wedge \diamond (Q \wedge \varphi_{i+2, i+3})) \right. \right. \\ \left. \left. \left. \vee \left(\diamond_{>1} (Q \wedge \varphi_{i+1, i+2}) \wedge \diamond_{<2} (R \wedge \varphi_{i+2, i+3}) \right) \right) \right) \right).$$

The same idea can be generalised to handle counting modalities and their past counterparts.

Proposition 3. $\text{MTL}[\mathcal{U}, \mathcal{S}] \stackrel{\leftrightarrow}{=} \text{MTL}[\mathcal{U}, \mathcal{S}, \{C_n, \overleftarrow{C}_n\}_{n=2}^\infty]$ over $[0, N)$ -timed words.

3.3 Non-Local Properties

The results in the last subsection shows that part of the expressiveness hierarchy collapses in the time-bounded setting. Nonetheless, $\text{MTL}[\mathcal{U}, \mathcal{S}]$ is still not expressive enough to capture all of $\text{FO}[<, +1]$. Recall that another feature of the continuous semantics used in the proof in [9] is that $\diamond_{=i}\varphi$ holds at t iff φ holds at $t+i$. Suppose that we want to specify the following property over $\mathbf{P} = \{P, Q\}$ at the current time t_1 for some integer constant $a > 0$:

- There is an event in $(t_1 + a, \infty)$ where Q holds
- P holds at all events between $t_1 + a$ and there.

Under the continuous semantics, by introducing a special monadic predicate P_ϵ that holds at all ‘no-event’ points in the flow, the property can easily be expressed as

$$\varphi_{cont1} = \diamond_{=a}((P \vee P_\epsilon) \mathcal{U} Q).$$

See Figure 3 for an illustration. Filled boxes denote events at which $\neg P \wedge Q$ holds whereas hollow boxes denote events at which $P \wedge \neg Q$ holds. The formula φ_{cont1} holds at t_1 in the continuous semantics. By contrast, it does not hold at t_1 in the pointwise semantics since there is no event at $t_1 + a$.




Fig. 3. φ_{cont1} holds at t_1 in the continuous semantics

To express the above property in the pointwise semantics, we introduce a relatively simple family of modalities $\mathcal{B}_I^\rightarrow$ (called ‘Beginning’) and their past versions \mathcal{B}_I^\leftarrow . They can be used to specify the *first* events in given intervals. For example, the following modality asserts that the first event in (a, b) satisfies φ :

$$\begin{aligned} \mathcal{B}_{(a,b)}^\rightarrow(x_1, \varphi) &= \exists x_2 (x_1 < x_2 \wedge d(x_1, x_2) > a \wedge d(x_1, x_2) < b \wedge \varphi(x_2) \\ &\quad \wedge \nexists x_3 (x_1 < x_3 \wedge x_3 < x_2 \wedge d(x_1, x_3) > a \wedge d(x_1, x_3) < b)). \end{aligned}$$

Now the property above can be defined as $\mathcal{B}_{(a,\infty)}^\rightarrow(P \mathcal{U}^w Q)$. We refer to the extension of $\text{MTL}[\mathcal{U}, \mathcal{S}]$ with $\mathcal{B}_I^\rightarrow, \mathcal{B}_I^\leftarrow$ as $\text{MTL}[\mathcal{U}, \mathcal{S}, \mathcal{B}^{\leftrightarrow}]$.

Proposition 4. $\text{MTL}[\mathcal{U}, \mathcal{S}] \subsetneq \text{MTL}[\mathcal{U}, \mathcal{S}, \mathcal{B}^{\leftrightarrow}]$ over $[0, N)$ -timed words.

Proof. This is a variation of the proof in [12, Section 5]. Given $m \in \mathbb{N}_{\geq 0}$, we describe models \mathcal{E}_m and \mathcal{F}_m that are indistinguishable by $\text{MTL}[\mathcal{U}, \mathcal{S}]$ formulas of modal depth $\leq m$ but distinguishing in $\text{MTL}[\mathcal{U}, \mathcal{S}, \mathcal{B}^{\leftrightarrow}]$. The detailed proof is given in Appendix.

We first describe \mathcal{E}_m . Let $g = \frac{1}{2m+6}$ and pick $\epsilon < \frac{g}{\frac{1}{g}-1}$. The first event (at time 0) satisfies $\neg P \wedge \neg Q$. Then, a sequence of overlapping segments starts at time $\frac{0.5}{2m+5}$ (see Figure 4). Each segment consists of an event satisfying $P \wedge \neg Q$ and an event satisfying $\neg P \wedge Q$. For ease of presentation we will refer to them as P -events and Q -events.




Fig. 4. A single segment in \mathcal{E}_m




Fig. 5. Near the $(3m+9)^{th}$ -segments in \mathcal{E}_m and \mathcal{F}_m

If the P -event in the i^{th} segment is at time t , then its Q -event is at time $t + 2g + \frac{1}{2} \cdot \epsilon$. All P -events in neighbouring segments are separated by $g - \frac{g}{\frac{1}{g}-1}$. We put a total of $4m+12$ segments. The only difference between the constructed models $\mathcal{E}_m, \mathcal{F}_m$ lies in the $(3m+9)^{th}$ segment. Let this segment starts at t_{3m+9} . In \mathcal{F}_m , we move the corresponding Q -event to $t + 2g - \frac{1}{2} \cdot \epsilon$ (see Figure 5). Note that there are P -events at time 0.5 in both models (in their $(m+4)^{th}$ segments). We have $\mathcal{E}_m \models \diamond(P \wedge \mathcal{B}_{(1,2)}^{\rightarrow} P)$ whilst $\mathcal{F}_m \not\models \diamond(P \wedge \mathcal{B}_{(1,2)}^{\rightarrow} P)$. □

Despite being a non-trivial addition to $\text{MTL}[\mathcal{U}, \mathcal{S}]$, we show that this simple extension is still less expressive than $\text{FO}[\langle, +1]$.

Proposition 5. $\text{MTL}[\mathcal{U}, \mathcal{S}, \mathcal{B}^{\leftrightarrow}] \subsetneq \text{FO}[\langle, +1]$ over $[0, N]$ -timed words.

Proof. Given $m \in \mathbb{N}_{\geq 0}$, we construct two models as follows. This is a variation of the proof in [13, Section 7]. Let

$$\mathcal{G}_m = (\emptyset, 0)(\emptyset, \frac{0.5}{2m+3})(\emptyset, \frac{1.5}{2m+3}) \dots (\emptyset, 1 - \frac{0.5}{2m+3}) \\ (\emptyset, 1 + \frac{0.5}{2m+2})(\emptyset, 1 + \frac{1.5}{2m+2}) \dots (\emptyset, 2 - \frac{0.5}{2m+2}).$$

\mathcal{H}_m is constructed as \mathcal{G}_m except that the event at time $\frac{m+1.5}{2m+3}$ is missing.




Fig. 6. Models \mathcal{G}_m and \mathcal{H}_m for $m = 2$

Figure 6 illustrates the models for the case $m = 2$. Filled boxes represent events. It can be shown that no $\text{MTL}[\mathcal{U}, \mathcal{S}, \mathcal{B}^{\text{tr}}]$ formula of modal depth $\leq m$ distinguishes \mathcal{G}_m and \mathcal{H}_m while there is a simple $\text{FO}[<, +1]$ formula does. The detailed proof is given in Appendix. \square

As the models \mathcal{G}_m and \mathcal{H}_m can, of course, be distinguished by an $\text{MTL}[\mathcal{U}, \mathcal{S}]$ formula in the continuous semantics, some further extension is needed in the pointwise framework. Observe that one can write formulas such as

$$\varphi_{cont2} = (\diamond_{=c}(P \vee P_\epsilon)) \mathcal{U} (\diamond_{=a}Q)$$

where a and c are distinct. This formula holds at t_1 (in the continuous semantics) iff there is $t' > t_1$ such that:

- There is an event at $t' + a$ and Q holds there
- P holds at all events in $(t_1 + c, t' + c)$.

See again Figure 3 for an illustration (φ_{cont2} also holds at t_1). This ability appears to be essential: let $\varphi' = \neg P_\epsilon \wedge \neg \diamond_{=1} \neg P_\epsilon$, the following formula distinguishes \mathcal{G}_m and \mathcal{H}_m in the continuous semantics:

$$\varphi_{cont-d} = \diamond (\diamond (\diamond_{=1} \neg P_\epsilon) \wedge \varphi' \wedge (P_\epsilon \wedge \diamond_{=1} P_\epsilon) \mathcal{U} (\varphi' \wedge \diamond \diamond_{=1} \neg P_\epsilon)).$$

In the next section, we propose new modalities that adds this ability to $\text{MTL}[\mathcal{U}, \mathcal{S}]$. We show later that this ability is indeed the ‘missing piece’ in our setting.

4 New Modalities

We introduce new unary modalities \mathfrak{E}_I^c (called ‘Earlier’), where $I \subseteq (0, \infty)$ is an interval with endpoints in $\mathbb{N}_{\geq 0} \cup \{\infty\}$, $c \in \mathbb{N}_{\geq 0}$ and $0 \leq c \leq l(I)$. In words, the temporal formula $\mathfrak{E}_{(a,b)}^c \varphi$ holds at t_1 iff the following holds:

- φ holds at the first event (at time t_2) in $(t_1 + a, t_1 + b)$
- There is no event in $(t_1 + c, t_1 + c + (t_2 - a))$.

For example, $\mathfrak{E}_{(a,\infty)}^c(P \wedge \neg Q)$ holds at t_1 in Figure 3 as $d_2 \leq d_1$. Formally,

$$\begin{aligned} \mathfrak{E}_{(a,b)}^c(x_1, \varphi) = & \exists x_2 \left(x_1 < x_2 \wedge d(x_1, x_2) > a \wedge d(x_1, x_2) < b \wedge \varphi(x_2) \right. \\ & \wedge \nexists x_3 (x_1 < x_3 \wedge x_3 < x_2 \wedge d(x_1, x_3) > a \wedge d(x_1, x_3) < b) \\ & \left. \wedge \nexists x_4 (x_1 < x_4 \wedge x_4 < x_2 \wedge d(x_1, x_4) > c \wedge d(x_2, x_4) > (a - c)) \right). \end{aligned}$$

Definitions for other types of intervals are similar. We also define their past counterparts \mathfrak{N}_I^c (called ‘Newer’). The logic obtained by adding these modalities to $\text{MTL}[\mathcal{U}, \mathcal{S}]$ is referred to as $\text{MTL}[\mathcal{U}, \mathcal{S}, \mathfrak{E}, \mathfrak{N}]$.

Our main result is the following theorem.

Theorem 1. *$\text{MTL}[\mathcal{U}, \mathcal{S}, \mathfrak{E}, \mathfrak{N}]$ is expressively complete for $\text{FO}[\langle, +1]$ over $[0, N]$ -timed words.*

4.1 More Liberal Bounds

It is easy to see that we can allow $c > l(I)$ in \mathfrak{E}_I^c for free. For instance, $\mathfrak{E}_{(2,5)}^{10} \varphi$ is equivalent to $\mathfrak{E}_{(2,5)}^2 \varphi \wedge \neg \mathfrak{E}_{(10,13)}^2 \mathbf{true}$ over $[0, N]$ -timed words. Indeed, we can even allow negative constants, i.e., $I \subseteq (-\infty, \infty)$ with endpoints in $\mathbb{Z} \cup \{\infty\}$ and $c \in \mathbb{Z}$. For example, over $[0, N]$ -timed words, $\mathfrak{E}_{(5,10)}^{-7} \varphi$ is equivalent to

$$\mathfrak{E}_{(5,10)}^5(\varphi \wedge \mathfrak{N}_{(5,10)}^{12} \mathbf{true}) \wedge \left(\mathfrak{E}_{(5,10)}^0 \mathbf{true} \vee \left(\varphi' \mathcal{U} (\varphi' \wedge \mathfrak{E}_{(5,10)}^0 \mathbf{true}) \right) \right)$$

where $\varphi' = \mathfrak{N}_{(0,\infty)}^7 \mathbf{true} \wedge \neg \diamond_{=7} \mathbf{true}$. Now let $\varphi'' = \mathbf{true} \wedge \neg \diamond_{=1} \mathbf{true}$. The models \mathcal{G}_m and \mathcal{H}_m in Section 3.3 can be distinguished by the following formula:

$$\varphi_{pw-d} = \diamond \left(\mathfrak{N}_{(1,\infty)}^1 \diamond^w \mathbf{true} \wedge \varphi'' \wedge \mathfrak{E}_{(0,\infty)}^{-1} (\varphi'' \wedge \mathfrak{E}_{(-1,\infty)}^{-1} \diamond^w \mathbf{true}) \right).$$

Compare this with the formula φ_{cont-d} defined in Section 3.3, which distinguishes \mathcal{G}_m and \mathcal{H}_m in the continuous semantics.

4.2 Generalised ‘Until’ and ‘Since’

To facilitate the proof of Theorem 1, we introduce a family of derived modalities which can be understood as generalisations of the usual ‘Until’ and ‘Since’

modalities. Let $I = (a, b) \subseteq (0, \infty)$ and $c \leq a$,

$$\begin{aligned} \mathfrak{U}_{(a,b)}^c(x_1, \varphi_1, \varphi_2) = & \exists x_2 \left(x_1 < x_2 \wedge d(x_1, x_2) > a \wedge d(x_1, x_2) < b \wedge \varphi_2(x_2) \right. \\ & \wedge \forall x_3 \left(x_1 < x_3 \wedge x_3 < x_2 \wedge d(x_1, x_3) > c \right. \\ & \left. \left. \wedge d(x_2, x_3) > (a - c) \implies \varphi_1(x_3) \right) \right). \end{aligned}$$

Definitions for other types of intervals are similar. We also define the generalised ‘Since’ modalities \mathfrak{S}_I^t .

Proposition 6. $\varphi_1 \mathfrak{U}_{(a,b)}^c \varphi_2$ with $(a, b) \subseteq (0, \infty)$ is definable in $\text{MTL}[\mathcal{U}, \mathcal{S}, \mathfrak{E}, \mathfrak{N}]$.

Proof. We give equivalent formulas (over $[0, N]$ -timed words) in each case below:

– $c = a$: It is clear that

$$\varphi_1 \mathfrak{U}_{(a,b)}^c \varphi_2 = \diamond_{(a,b)} \varphi_2 \wedge \mathfrak{E}_{(a,b)}^a(\varphi_1 \mathcal{U}^w \varphi_2).$$

– $c < a$: Let

$$\begin{aligned} \varphi_3 = \varphi_2 \vee & \left(\mathbf{true} \wedge \neg \diamond_{=(a-c)} \mathbf{true} \right. \\ & \left. \wedge \left(\neg \diamond_{=(a-c)} \mathbf{true} \wedge \mathfrak{N}_{(0,\infty)}^{(a-c)} \mathbf{true} \right) \mathcal{U} \left(\varphi_2 \wedge \mathfrak{N}_{(0,\infty)}^{(a-c)} \mathbf{true} \right) \right). \end{aligned}$$

The desired formula is the conjunction of $\diamond_{(a,b)} \varphi_2$ and

$$\mathfrak{E}_{(a,b)}^c \varphi_3 \vee \mathfrak{E}_{(c, c+(b-a))}^c \left(\varphi_1 \mathcal{U}^w \left(\varphi_1 \wedge \mathfrak{E}_{(a-c, \infty)}^0 \varphi_3 \right) \right).$$

– $c > a$: Let

$$\begin{aligned} \varphi_4 = \varphi_2 \vee & \left(\mathbf{true} \wedge \neg \diamond_{=(c-a)} \mathbf{true} \right. \\ & \left. \wedge \left(\neg \diamond_{=(c-a)} \mathbf{true} \wedge \mathfrak{N}_{(0,\infty)}^{(a-c)} \mathbf{true} \right) \mathcal{U} \left(\varphi_2 \wedge \mathfrak{N}_{(0,\infty)}^{(a-c)} \mathbf{true} \right) \right). \end{aligned}$$

The desired formula is the conjunction of $\diamond_{(a,b)} \varphi_2$ and

$$\mathfrak{E}_{(a,b)}^c \varphi_4 \vee \mathfrak{E}_{(c, c+(b-a))}^c \left(\varphi_1 \mathcal{U}^w \left(\varphi_1 \wedge \mathfrak{E}_{(a-c, \infty)}^0 \varphi_4 \right) \right).$$

□

The proofs for other type of intervals (and more liberal bounds) are similar. Note that the usual modalities can be written in the generalised modalities, e.g.,

$$\varphi_1 \mathcal{U}_{(a,b)} \varphi_2 = \varphi_1 \mathfrak{U}_{(a,b)}^a \varphi_2 \wedge \neg \left(\mathbf{true} \mathfrak{U}_{(0,a]}^0 (\neg \varphi_1) \right).$$

We will refer to the logic with modalities $\mathfrak{U}_I^c, \mathfrak{S}_I^c$ (where $I \subseteq (-\infty, \infty)$ with endpoints in $\mathbb{Z} \cup \{\infty\}$ and $c \in \mathbb{Z}$) as $\text{MTL}[\mathfrak{U}, \mathfrak{S}]$.

Proposition 7. $\text{MTL}[\mathfrak{U}, \mathfrak{S}] \stackrel{\dagger}{=} \text{MTL}[\mathcal{U}, \mathcal{S}, \mathfrak{E}, \mathfrak{N}]$ over \mathbb{T} -timed words.

5 The Translation

We give a translation from an arbitrary $\text{FO}[<, +1]$ formula with one free variable into an equivalent $\text{MTL}[\mathfrak{U}, \mathfrak{S}]$ formula. Our proof strategy is similar to that in [9]: we eliminate the metric by introducing extra predicates, convert to LTL , and then replace the new predicates by their equivalent $\text{MTL}[\mathfrak{U}, \mathfrak{S}]$ formulas.

5.1 Eliminating the Metric

We introduce fresh monadic predicates $\overline{\mathbf{P}} = \{P_i \mid P \in \mathbf{P}, 0 \leq i \leq N-1\}$ as in [9] and, additionally, $\overline{\mathbf{Q}} = \{Q_i \mid 0 \leq i \leq N-1\}$. The monadic predicates $\overline{\mathbf{Q}}$ indicate the existence of events. Let $\varphi_{event} = \forall x \left(\bigvee_{i \in [0, N-1]} Q_i(x) \right) \wedge \forall x \left(\bigwedge_{i \in [0, N-1]} (P_i(x) \implies Q_i(x)) \right)$ and $\varphi_{init} = \exists x (\#x' (x' < x) \wedge Q_0(x))$. There is an obvious ‘stacking’ bijection (indicated by overlining) between $[0, N]$ -timed words over $2^{\mathbf{P}}$ and $[0, 1]$ -timed words over $2^{\overline{\mathbf{P}} \cup \overline{\mathbf{Q}}}$ satisfying $\varphi_{event} \wedge \varphi_{init}$.

Let $\vartheta(x)$ be an $\text{FO}[<, +1]$ formula with one free variable and in which each quantifier uses a fresh new variable. Wlog. we assume that $\vartheta(x)$ contains only existential quantifiers (this can be achieved by syntactic rewriting). Replace the formula by

$$(Q_0(x) \wedge \vartheta[x/x]) \vee (Q_1(x) \wedge \vartheta[x+1/x]) \vee \dots \vee (Q_{N-1}(x) \wedge \vartheta[x+(N-1)/x])$$

where $\vartheta[e/x]$ denotes the formula obtained by substituting all free occurrences of x in ϑ by (an expression) e . Then, similarly, recursively replace every subformula $\exists x' \theta$ by

$$\exists x' ((Q_0(x') \wedge \theta[x'/x']) \vee \dots \vee (Q_{N-1}(x') \wedge \theta[x'+(N-1)/x'])).$$

Note that we do not actually have the $+1$ function in our structures; they only serve as annotations and will be removed later, e.g., $x' + k$ means that $Q_k(x')$ holds. Then, carry out the following syntactic substitutions:

- For each inequality of the form $x_1 + k_1 < x_2 + k_2$, replace it with
 - $x_1 < x_2$ if $k_1 = k_2$
 - **true** if $k_1 < k_2$
 - **¬true** if $k_1 > k_2$
- For each distance formula, e.g., $d(x_1 + k_1, x_2 + k_2) \leq 2$, replace it with
 - **true** if $|k_1 - k_2| \leq 1$
 - $(\neg(x_1 < x_2) \wedge \neg(x_2 < x_1)) \vee (x_2 < x_1)$ if $k_2 - k_1 = 2$
 - $(\neg(x_1 < x_2) \wedge \neg(x_2 < x_1)) \vee (x_1 < x_2)$ if $k_1 - k_2 = 2$
 - **¬true** if $|k_1 - k_2| > 2$
- Replace terms of the form $P(x_1 + k)$ with $P_k(x_1)$.

This gives a non-metric first-order formula $\overline{\vartheta}(x)$ over $\overline{\mathbf{P}} \cup \overline{\mathbf{Q}}$. Denote by $frac(t)$ the fractional part of a non-negative real t . It is not hard to see that for each $[0, N]$ -timed word $\rho = (\sigma, \tau)$ over $2^{\mathbf{P}}$ and its stacked counterpart $\overline{\rho}$, the following holds:

- $\rho, t \models \vartheta(x)$ implies $\bar{\rho}, \bar{t} \models \bar{\vartheta}(x)$ where $\bar{t} = \text{frac}(t)$
- $\bar{\rho}, \bar{t} \models \bar{\vartheta}(x)$ implies there exists $t \in U_\rho$ with $\text{frac}(t) = \bar{t}$ s.t. $\rho, t \models \vartheta(x)$.

Moreover, if $\rho, t \models \vartheta(x)$, then the integral part of t indicates which clause in $\bar{\vartheta}(x)$ is satisfied when x is substituted with $\bar{t} = \text{frac}(t)$, and vice versa.

By Kamp's theorem [7], $\bar{\vartheta}(x)$ is equivalent to an LTL $[\mathcal{U}, \mathcal{S}]$ formula $\bar{\varphi}$ of the following form:

$$(Q_0 \wedge \bar{\varphi}_0) \vee (Q_1 \wedge \bar{\varphi}_1) \vee \dots \vee (Q_{N-1} \wedge \bar{\varphi}_{N-1}).$$

5.2 From Non-Metric to Metric

We now construct an MTL $[\mathcal{U}, \mathcal{S}]$ formula that is equivalent to $\vartheta(x)$ over $[0, N)$ -timed words. Note that we make heavy use of the formulas in \mathcal{F}_{unit} defined in Section 3.1.

Proposition 8. *Let $\bar{\psi}$ be a subformula of $\bar{\varphi}_i$ for some $i \in [0, N-1)$. There is an MTL $[\mathcal{U}, \mathcal{S}]$ formula ψ such that for all $\bar{t} \in \bar{\rho}$ and $t \in \rho$ with $\text{frac}(t) = \bar{t}$, we have*

$$\bar{\rho}, \bar{t} \models \bar{\psi} \iff \rho, t \models \psi.$$

Proof. We prove this by induction on the structure of $\bar{\varphi}_i$.

– *Base step.* Consider the following cases:

- $\bar{\psi} = P_j$: Let

$$\psi = (\varphi_{0,1} \wedge \diamond_{=j} P) \vee \dots \vee (\varphi_{j,j+1} \wedge P) \vee \dots \vee (\varphi_{N-1,N} \wedge \diamond_{=((N-1)-j)} P).$$

- $\bar{\psi} = Q_j$: Similarly we let

$$\psi = (\varphi_{0,1} \wedge \diamond_{=j} \mathbf{true}) \vee \dots \vee (\varphi_{j,j+1} \wedge \mathbf{true}) \vee \dots \vee (\varphi_{N-1,N} \wedge \diamond_{=((N-1)-j)} \mathbf{true}).$$

– *Induction step.* The case for boolean operations are trivial and hence omitted.

- $\bar{\psi} = \bar{\psi}_1 \mathcal{U} \bar{\psi}_2$: By IH we have ψ_1 and ψ_2 . Let

$$\psi^{j,k,l} = \psi_1 \mathcal{U}_{(j,j+1)}^k (\psi_2 \wedge \varphi_{l,l+1}).$$

The desired formula is

$$\psi = \bigvee_{i \in [0, N-1]} \left(\varphi_{i,i+1} \wedge \bigvee_{\substack{j \in [-i, \dots, (N-1)-i] \\ l = i+j}} \left(\bigwedge_{k \in [-i, \dots, (N-1)-i]} \psi^{j,k,l} \right) \right).$$

- $\bar{\psi} = \bar{\psi}_1 \mathcal{S} \bar{\psi}_2$: This is symmetric to the case for $\bar{\psi}_1 \mathcal{U} \bar{\psi}_2$. □

Construct corresponding formulas φ_i for each $\bar{\varphi}_i$ using the proposition above. Substitute them into $\bar{\varphi}$ and replace all Q_i by $\varphi_{i,i+1}$ to obtain our final formula φ . We claim that it is equivalent to $\vartheta(x)$ over $[0, N)$ -timed words.

Proposition 9. *For all \mathbb{T} -timed words ρ and $t \in U_\rho$, we have*

$$\rho, t \models \varphi(x) \iff \rho, t \models \vartheta(x).$$

Proof. Follows directly from Section 5.1 and Proposition 8. □

References

1. Bouyer, P., Chevalier, F., Markey, N.: On the expressiveness of TPTL and MTL. In: Proceedings of FSTTCS 2005. LNCS, vol. 3821, pp. 432–443. Springer (2005)
2. D’Souza, D., Prabhakar, P.: On the expressiveness of MTL in the pointwise and continuous semantics. *International Journal on Software Tools for Technology Transfer* 9(1), 1–4 (2006)
3. Hirshfeld, Y., Rabinovich, A.: Expressiveness of metric modalities for continuous time. In: Proceedings of CSR 2006. LNCS, vol. 3967, pp. 211–220. Springer (2006)
4. Hirshfeld, Y., Rabinovich, A.: Expressiveness of metric modalities for continuous time. *Logical Methods in Computer Science* 3(1) (2007)
5. Hunter, P.: When is metric temporal logic expressively complete? In: Proceedings of CSL 2013. LIPIcs, vol. 23, pp. 380–394. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik (2013)
6. 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)
7. Kamp, J.: Tense logic and the theory of linear order. Ph.D. thesis, University of California, Los Angeles (1968)
8. Koymans, R.: Specifying real-time properties with metric temporal logic. *Real-Time Systems* 2(4), 255–299 (1990)
9. Ouaknine, J., Rabinovich, A., Worrell, J.: Time-bounded verification. In: Proceedings of CONCUR 2009. LNCS, vol. 5710, pp. 496–510. Springer (2009)
10. Ouaknine, J., Worrell, J.: On the decidability of metric temporal logic. In: Proceedings of LICS 2005. pp. 188–197. IEEE Computer Society Press (2005)
11. Ouaknine, J., Worrell, J.: Towards a theory of time-bounded verification. In: Proceedings of ICALP 2010. LNCS, vol. 6199, pp. 22–37. Springer (2010)
12. 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)
13. Prabhakar, P., D’Souza, D.: On the expressiveness of MTL with past operators. In: Proceedings of FORMATS 2006. LNCS, vol. 4202, pp. 322–336. Springer (2006)

A MTL[$\mathcal{U}, \mathcal{S}, \mathcal{B}^{\leftrightarrow}$] EF Games

We resort to a variant of the MTL[\mathcal{U}, \mathcal{S}] EF theorem for the following proofs. A brief account of the underlying EF games is outlined below. For details we refer the reader to [12].

An m -round MTL[$\mathcal{U}, \mathcal{S}, \mathcal{B}^{\leftrightarrow}$] EF game starts with round 0 and ends with round m . A *configuration* is a pair of positions (i, i') , respectively in two timed words. Let (i_r, j_r) be the configuration at the beginning of round r . *Spoiler* first checks both events satisfy the same set of monadic predicates. Then she chooses (say) i'_r according to I . *Duplicator* chooses j'_r as a response, also according to I . If *Spoiler* plays \diamondleftarrow -part, \diamondrightarrow -part, or it is indeed a $\mathcal{B}^{\rightarrow}$ -move or \mathcal{B}^{\leftarrow} -move,³ the game proceeds to the next round with $(i_{r+1}, j_{r+1}) = (i'_r, j'_r)$. If she plays \mathcal{U} -part or \mathcal{S} -part, another position j''_r would be chosen by her, and *Duplicator* would need to choose i''_r as response. The game then proceeds to the next round with $(i_{r+1}, j_{r+1}) = (i''_r, j''_r)$. If *Duplicator* fails to respond at any point then *Spoiler* wins the game. We write $\rho, i \approx_m \rho', i'$ if *Duplicator* has a winning strategy for the m -round EF game on ρ, ρ' that starts from configuration (i, i') .⁴ In the theorem below, modalities allowed in φ correspond to moves that *Spoiler* can take in the game.

Theorem 2 ([12]). *For timed words ρ, ρ' and a temporal formula φ of modal depth $\leq m$,*

$$\rho, 1 \approx_m \rho', 1 \text{ implies } \rho \models \varphi \iff \rho' \models \varphi.$$

B Proof of Proposition 1

Proof. We play the game on $\mathcal{A}_m, \mathcal{B}_m$. It is clear that if $i_r = j_r$ and $i_r \geq 2$ then *Duplicator* wins the remaining rounds. If $i_r = j_r + 1$ and *Spoiler* chooses some move, *Duplicator* can always make $i_{r+1} = j_{r+1} \geq 2$ or $(i_{r+1}, j_{r+1}) = (i_r + 1, j_r + 1)$. The result follows from Theorem 2. \square

C Proof of Proposition 2

Proof. We play the game on $\mathcal{C}_m, \mathcal{D}_m$. If $i_r = j_r + 1$ and $i_r \geq m + 4$ then *Duplicator* wins the remaining rounds. The proof is similar to the proof of Proposition 2. \square

D Proof of Proposition 4

As noted before, the only difference in two models is a pair of Q -events. We denote this pair of events by x and y respectively and write their corresponding timestamps as t_x and t_y (see Figure 5). It is easy to verify that no two events

³ We define additional moves for $\mathcal{B}_I^{\rightarrow}$ and $\mathcal{B}_I^{\leftarrow}$ according to their semantics.

⁴ For our purpose, the largest constant used in the games can be assumed to be ≤ 2 in all proofs.

are separated by an integer distance. We say a configuration (i, j) is *identical* if $i = j$. For $i \geq 2$, we denote by $seg(i)$ the segment that the i^{th} event belongs to, and we write $P(i)$ if the i^{th} event is a P -event and $Q(i)$ if its a Q -event.

Proposition 10. *Duplicator has a winning strategy for m -round $MTL[\mathcal{U}, \mathcal{S}]$ EF game on $\mathcal{E}_m, \mathcal{F}_m$ that starts from $(1, 1)$. In particular, she has a winning strategy such that for each round $0 \leq r \leq m$, the i_r^{th} event in \mathcal{E}_m and the j_r^{th} event in \mathcal{F}_m satisfy the same set of propositions and*

- if $i_r \neq j_r$, then
 - $seg(i_r) - seg(j_r) < r$
 - $(m + 1 - r) < seg(i_r), seg(j_r) < (m + 5 + r)$ or $(3m + 8 - r) < seg(i_r), seg(j_r) < (3m + 12 + r)$.

Proof. We prove the proposition by induction on r .

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

- $(i_r, j_r) = (1, 1)$:
Duplicator can always make (i_{r+1}, j_{r+1}) identical.
- $(i_r, j_r) \neq (1, 1)$ is identical:
Duplicator tries to make (i'_r, j'_r) identical. This may only fail when
 - * $P(i_r) \wedge P(j_r)$ and $seg(i_r) = seg(j_r) = m + 4$.
 - * $Q(i_r) \wedge Q(j_r)$ and $seg(i_r) = seg(j_r) = 3m + 9$, i.e., x and y .

In these cases, *Duplicator* chooses another event in a neighbouring segment that minimises $|seg(i'_r) - seg(j'_r)|$. For example, if (i_r, j_r) corresponds to x and y and *Spoiler* chooses j'_r such that $P(j'_r)$ and $seg(j'_r) = m + 4$ in a $\mathcal{S}_{(1, \infty)}$ -move, *Duplicator* chooses i'_r with $seg(i'_r) = m + 3$. If *Spoiler* then plays \Leftarrow -part, the resulting configuration (i_{r+1}, j_{r+1}) will clearly satisfy the claim. If she plays \mathcal{S} -part, *Duplicator* makes (i''_r, j''_r) identical whenever possible. Otherwise she chooses the appropriate event that minimises $|seg(i''_r) - seg(j''_r)|$. For instance, if $Q(i''_r)$ and $seg(i''_r) = m + 1$, *Duplicator* chooses j''_r such that $Q(j''_r)$ and $seg(j''_r) = m + 2$.

- (i_r, j_r) is not identical:
Duplicator tries to make (i'_r, j'_r) identical. If this is not possible, then *Duplicator* chooses an event that minimises $|seg(i'_r) - seg(j'_r)|$. For example, consider $seg(i_r) = m + 4$, $seg(j_r) = m + 3$ such that $P(i_r)$ and $P(j_r)$, and *Spoiler* chooses x in an $\mathcal{U}_{(0, 1)}$ -move. In this case, *Duplicator* cannot choose y' but the first Q -event that happens before y' . *Duplicator* responds to \mathcal{U} -parts and \mathcal{S} -parts in similar ways as before. It is easy to see that the claim holds. □

Proposition 4 now follows directly from Proposition 10.

E Proof of Proposition 5

Observe that no two events are separated by an integer distance. We say a configuration (i, j) is *synchronised* if they correspond to events with the same timestamp.

Proposition 11. *Duplicator has a winning strategy for m -round $\text{MTL}[\mathcal{U}, \mathcal{S}, \mathcal{B}^{\leftrightarrow}]$ EF game on $\mathcal{G}_m, \mathcal{H}_m$ that starts from $(1, 1)$. In particular, she has a winning strategy such that for each round $0 \leq r \leq m$, the i_r^{th} event in \mathcal{G}_m and the j_r^{th} event in \mathcal{H}_m satisfy the same set of propositions and*

- if (i_r, j_r) is not synchronised, then
 - $|i_r - j_r| = 1$
 - $(m + 2 - r) < i_r, j_r < (m + 4 + r)$ or $(3m + 5 - r) < i_r, j_r < (3m + 6 + r)$.

Proof. We prove the proposition by induction on r .

- *Base step.* The proposition holds trivially for $(i_0, j_0) = (1, 1)$.
- *Induction step.* Suppose that the claim holds for $r < m$. We prove it also holds for $r + 1$.
 - $(i_r, j_r) = (1, 1)$:
Duplicator tries to make (i'_r, j'_r) synchronised. If *Spoiler* chooses $i'_r = m + 3$, *Duplicator* chooses $j'_r = m + 2$.
 - $(i_r, j_r) \neq (1, 1)$ is synchronised:
Duplicator tries to make (i'_r, j'_r) synchronised. If this is not possible then *Duplicator* chooses the event that minimises $|i'_r - j'_r|$. It is easy to see that the resulting configuration (i_{r+1}, j_{r+1}) satisfies the claim regardless of how *Spoiler* plays.
 - (i_r, j_r) is not synchronised:
 The strategy of *Duplicator* is same as the case above.

□

Proposition 5 now follows directly from Proposition 11.