On the Expressiveness of Metric Temporal Logic over Bounded Timed Words

Hsi-Ming Ho

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 ($\mathsf{FO}[<, +1]$) in the pointwise semantics over bounded time domains (i.e., timed words of bounded duration) [20]. In this paper, we present an extension of MTL which has the same expressive power as $\mathsf{FO}[<, +1]$ 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 [14] states that, in either case, LTL has precisely the same expressive power as the *Monadic First-Order Logic of Order* ($\mathsf{FO}[<]$). 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 [15] 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*. Timed words are the leading interpretation, for example, for systems modelled as timed automata [1]. 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., [17] for a historical account.

Alongside these developments, researchers proposed the *Monadic First-Order Logic of Order and Metric* ($\mathsf{FO}[<, +1]$) as a natural quantitative extension of $\mathsf{FO}[<]$. Unfortunately, Hirshfeld and Rabinovich [10] showed that no 'finitary' extension of MTL—and *a fortiori* MTL itself—could have the same expressive power as $\mathsf{FO}[<, +1]$ over the reals.[1] Still, in the continuous semantics, MTL can be made expressively complete for $\mathsf{FO}[<, +1]$ by extending the logic with an infinite family of '*counting modalities*' [12] or considering only *bounded* time domains [16, 18]. Nonetheless, and rather surprisingly, MTL with counting modalities remains strictly less expressive than $\mathsf{FO}[<, +1]$ over bounded time domains in the pointwise semantics, i.e., over timed words of bounded duration, as we will see in Section 3.

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

It is worth noting that $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ satisfiability and model checking (against timed automata) are decidable over bounded time domains, thanks to the decidability of $\mathsf{FO}[<, +1]$ over such domains as established in [16, 18]. Unfortunately, $\mathsf{FO}[<, +1]$ has non-elementary complexity, whereas the time-bounded satisfiability and model-checking problems for MTL are EXPSPACE-complete [16, 18]. However, it can easily be seen by inspecting the relevant constructions that the complexity bounds for MTL carry over to our new logic $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$.

2 Preliminaries

2.1 Timed Automata

2.2 Timed Words

Let the time domain $\mathbb{T}$ be a subinterval of $\mathbb{R}_{\geq 0}$ that contains 0. A *time sequence* $\tau = \tau_0\tau_1 \ldots$ is a non-empty finite or infinite sequence over $\mathbb{T}$ (called *timestamps*) that satisfies the requirements below (we denote the length of τ by $|\tau|$):

- *Initialisation*[2]: $\tau_0 = 0$
- *Strict monotonicity*: For all i, $0 \leq i < |\tau| - 1$, we have $\tau_i < \tau_{i+1}$.

If τ is infinite we require it to be unbounded, i.e., we disallow so-called Zeno sequences. A ***timed word*** over finite alphabet Σ is a pair $\rho = (\sigma, \tau)$, where $\sigma = \sigma_0\sigma_1 \ldots$ is a non-empty finite or infinite word over Σ and τ is a time sequence over $\mathbb{T}$ of the same length. We refer the pair (σ_i, τ_i) as the i^{th} *event* in

[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 Prabhakar and D'Souza [20] and Pandya and Shah [19] independently showed that MTL is strictly weaker than $\mathsf{FO}[<, +1]$ in the pointwise semantics.

[2] This requirement is natural in the present context as all the logics we consider in this thesis are *translation invariant*: two timed words are indistinguishable by formulas (of these logics) if they only differ by a fixed delay.

ρ, and define the *distance* between i^{th} and j^{th} ($i \leq j$) events to be $\tau_j - \tau_i$. In this sense, a timed word can be equivalently 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 $0 \leq i < |\rho|$. The *duration* of ρ is defined as $\tau_{|\rho|-1}$ if ρ is finite. We write $t \in \rho$ if t is equal to one of the timestamps in ρ.

In the present paper, we are mainly concerned with *finite* timed words with *bounded* time domains of the form $[0, N)$ where $N \in \mathbb{N}$ (we call them *bounded timed words*). For clarity, we refer to finite timed words with all timestamps in $\mathbb{T}$ as $\mathbb{T}$-timed words.

2.3 Flows

A ***flow*** over finite alphabet Σ is a function $f : \mathbb{T} \mapsto \Sigma$ that is *finitely variable*, i.e., the restriction of f to a subinterval of $\mathbb{T}$ of finite length has only finite number of discontinuities. We sometimes write $\mathbb{T}$-flows for flows with time domain $\mathbb{T}$.

2.4 Metric Logics

Syntax We first define a metric predicate logic $\mathsf{FO}[<, +1]$, of which all other logics discussed in this thesis will be defined as sublogics. In the sequel, we write $\Sigma_{\mathbf{P}} = 2^{\mathbf{P}}$ for a set of monadic predicates $\mathbf{P}$.

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

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

where $P \in \mathbf{P}$*,* x, x' *are variables,* $\sim \in \{=, \neq, <, >, \leq, \geq\}$ *and* $c \in \mathbb{N}$.[3]

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

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

The MTL formula $\varphi_1\, \mathcal{U}_{(0,5)}\, \varphi_2$ (using infix notation) is obtained by substituting MTL formulas φ_1, φ_2 for X_1, X_2, respectively.

[3] Note that whilst we still refer to the logic as $\mathsf{FO}[<, +1]$, we adopt here an equivalent definition using a binary distance predicate $d(x, x')$ (as in [21]) in place of the usual $+1$ function symbol.

Definition 2. *Given a set of monadic predicates* $\mathbf{P}$*, the set of* MTL *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* $\mathbb{N} \cup \{\infty\}$.

The (future-only) fragment $\mathsf{MTL}_{\mathsf{fut}}$ is obtained by banning subformulas of the form $\varphi_1 \, \mathcal{S}_I \, \varphi_2$. If I is not present as a subscript to a given modality then it is assumed to be $(0, \infty)$. We sometimes use pseudo-arithmetic expressions to denote intervals, e.g., '≥ 1' denotes $[1, \infty)$ and '$= 1$' denotes the singleton $\{1\}$. We also employ the usual syntactic sugar, e.g., $\mathbf{false} \equiv \neg\mathbf{true}$, $\Diamond_I \varphi \equiv \mathbf{true} \, \mathcal{U}_I \, \varphi$, $\overleftarrow{\Diamond}_I \varphi \equiv \mathbf{true} \, \mathcal{S}_I \, \varphi$, $\Box_I \varphi \equiv \neg \Diamond_I \neg \varphi$ and $\bigcirc_I \varphi \equiv \mathbf{false} \, \mathcal{U}_I \, \varphi$, etc.

Pointwise Semantics With each timed word $\rho = (\sigma, \tau)$ over $\Sigma_{\mathbf{P}}$ we associate a structure M_ρ. Its universe U_ρ is the subset $\{\tau_i \mid 0 \leq i < |\rho|\}$ of $\mathbb{T}$. The order relation $<$ and monadic predicates in $\mathbf{P}$ are interpreted in the expected way. For example, $P(\tau_i)$ holds in M_ρ iff $P \in \sigma_i$. The binary *distance predicate* $d(x, x') \sim c$ holds iff $|x - x'| \sim c$. The satisfaction relation is defined inductively as usual. We write $M_\rho, t_0, \ldots, t_n \models \vartheta(x_0, \ldots, x_n)$ (or $\rho, t_0, \ldots, t_n \models \vartheta(x_0, \ldots, x_n)$) if $t_0, \ldots, t_n \in U_\rho$ and $\vartheta(t_0, \ldots, t_n)$ holds in M_ρ. We say that $\mathsf{FO}[<, +1]$ formulas $\vartheta_1(x)$ and $\vartheta_2(x)$ are *equivalent* over $\mathbb{T}$-timed words if for all $\mathbb{T}$-timed words ρ and $t \in U_\rho$,

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

We say a metric logic L' is ***expressively complete*** for metric logic L over $\mathbb{T}$-timed words iff for any formula $\vartheta(x) \in L$, there is an equivalent formula $\varphi(x) \in L'$ over $\mathbb{T}$-timed words.

As we have seen earlier, each MTL formula can be defined as an $\mathsf{FO}[<, +1]$ formula with a single free first-order variable. Here for the sake of completeness we give an (equivalent) traditional inductive definition of the satisfaction relation for MTL over timed words.

Definition 3. *The satisfaction relation* $(\rho, i) \models \varphi$ *for an* MTL *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, $i < j < |\rho|$ *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 temporal modalities, e.g., φ_2 holds at i does not imply that $\varphi_1 \, \mathcal{U} \, \varphi_2$ holds at i. We write $\rho \models \varphi$ if $(\rho, 0) \models \varphi$.

Continuous Semantics With each flow f over $\Sigma_{\mathbf{P}}$ we associate a structure M_f. Its universe U_f is $\mathbb{T}$. The order relation $<$ and monadic predicates in $\mathbf{P}$ are interpreted in the expected way, e.g., $P(x)$ holds in M_f iff $P \in f(x)$. The binary *distance predicate* $d(x, x') \sim c$ holds iff $|x - x'| \sim c$. We write $M_f, t_0, \ldots, t_n \models \vartheta(x_0, \ldots, x_n)$ (or $f, t_0, \ldots, t_n \models \vartheta(x_0, \ldots, x_n)$) if $t_0, \ldots, t_n \in U_f$ and $\vartheta(t_0, \ldots, t_n)$ holds in M_f. We say that $\mathsf{FO}[<, +1]$ formulas $\vartheta_1(x)$ and $\vartheta_2(x)$ are equivalent over $\mathbb{T}$-flows if for all $\mathbb{T}$-flows f and $t \in U_f$,

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

A metric logic L' is ***expressively complete*** for metric logic L over $\mathbb{T}$-flows iff for any formula $\vartheta(x) \in L$, there is an equivalent formula $\varphi(x) \in L'$ over $\mathbb{T}$-flows.

The satisfaction relation for MTL over flows is defined as follows.

Definition 4. *The satisfaction relation $(f, t) \models \varphi$ for an MTL formula φ, a flow f and $t \in U_f$ is defined as follows:*

- $(f, t) \models P$ *iff* $P(t)$ *holds in* M_f
- $(f, t) \models \mathbf{true}$
- $(f, t) \models \varphi_1 \wedge \varphi_2$ *iff* $(f, t) \models \varphi_1$ *and* $(f, t) \models \varphi_2$
- $(f, t) \models \neg\varphi$ *iff* $(f, t) \not\models \varphi$
- $(f, t) \models \varphi_1 \mathcal{U}_I \varphi_2$ *iff there exists* $t' > t$, $t' \in \mathbb{T}$ *such that* $(f, t') \models \varphi_2$, $t' - t \in I$ *and* $(f, t'') \models \varphi_1$ *for all* t'' *with* $t < t'' < t'$
- $(f, t) \models \varphi_1 \mathcal{S}_I \varphi_2$ *iff there exists* $t' < t$, $t' \in \mathbb{T}$ *such that* $(f, t') \models \varphi_2$, $t - t' \in I$ *and* $(f, t'') \models \varphi_1$ *for all* t'' *with* $t' < t'' < t$.

We write $f \models \varphi$ if $(f, 0) \models \varphi$.

Relating the two Semantics Note that timed words can be regarded as a particular kind of flow: for a given $\mathbb{T}$-timed word ρ over Σ_P, we can construct a corresponding $\mathbb{T}$-flow f^ρ over $\Sigma_{P'}$, where $P' = P \cup \{P_\epsilon\}$, as follows:

- $f^\rho(\tau_i) = \sigma_i$ for all i, $0 \leq i < |\rho|$
- $f^\rho(\tau_i) = \{P_\epsilon\}$.

This enables us to interpret metric logics over timed words 'continuously'. We can thus compare the expressiveness of metric logics in both semantics by restricting the models of the continuous interpretations of metric logics to flows of this form (i.e., f^ρ for some timed word ρ). For example, we say that $\mathsf{FO}[<, +1]$ is not less expressive in the continuous semantics than in the pointwise semantics since for each $\mathsf{FO}[<, +1]$ formula $\vartheta_{pw}(x)$, there is an 'equivalent' $\mathsf{FO}[<, +1]$ formula $\vartheta_{cont}(x)$ such that $\rho, t \models \vartheta_{pw}(x)$ iff $f^\rho, t \models \vartheta_{cont}(x)$.

It is worth pointing out that the range of first-order quantifiers is the most crucial difference between the pointwise and continuous interpretations of metric logics. For example, the $\mathsf{MTL}_{\mathsf{fut}}$ formula $\Diamond_{=1}\mathbf{true}$ does not hold at position 0 in $\rho = (\sigma_0, 0)(\sigma_1, 0.5)(\sigma_2, 1.2)$ since there is no event at time 1; but the same formula holds at time 0 in f^ρ. While the ability to quantify over time points between events appears to add to the expressiveness of metric logics, this is not the case for $\mathsf{FO}[<,+1]$ as both interpretations indeed have equal expressiveness (when one considers only flows of the form f^ρ) [6].[4] However, it is known that MTL is strictly more expressive in the continuous semantics than in the pointwise semantics [20]. Combined with the expressive equivalence of $\mathsf{FO}[<,+1]$ in both semantics, this implies that MTL is not expressive complete for $\mathsf{FO}[<,+1]$ over bounded timed words.

2.5 MTL EF Games

In the proofs of the expressiveness results in this paper, we resort to (extended versions of) the MTL EF theorem given in [19], which itself is a timed generalisation of the LTL EF theorem [8]. A brief account of the underlying EF games played on a pair of timed words is outlined below.

An m-round MTL EF game starts with round 0 and ends with round m. A *configuration* is a pair of positions (i,j), respectively in two timed words ρ,ρ'. 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 one of the two timed words and an interval I; as an example say that she chooses ρ. She then picks i'_r such that $\tau_{i'_r}-\tau_{i_r}\in I$, where $\tau_{i'_r}$ and τ_{i_r} are the corresponding timestamps in ρ. *Duplicator* must choose a position j'_r in ρ' such that the difference of the corresponding timestamps in ρ' is in I. If *Spoiler* plays $\diamondsuit$-part or $\overleftarrow{\diamondsuit}$-part, 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 in ρ' such that $j_r<j''_r<j'_r$ (if exists) would be chosen by her, and *Duplicator* would need to choose a position i''_r in ρ such that $i_r<i''_r<i'_r$ 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 $\rho,i\approx_m\rho',j$ if *Duplicator* has a winning strategy for the m-round MTL EF game on ρ,ρ' that starts from configuration (i,j).[5]

Theorem 1 ([19]). *For timed words ρ,ρ' and an MTL formula φ of modal depth $\leq m$,*

$$\rho,0\approx_m\rho',0 \textit{ implies } \rho\models\varphi \iff \rho'\models\varphi\,.$$

Note that in the theorem above, moves allowed in the game correspond to modalities allowed in φ. In the sequel, we will introduce new moves into the game based on newly introduced modalities. The theorem extends in a straightforward manner.

3 A Hierarchy of Expressiveness

In this section, we present a sequence of successively more expressive extensions of $\mathsf{MTL}_{\mathsf{fut}}$ over bounded timed words. Along the way we highlight the key fea-

[4] The translation in [6] also holds in a time-bounded setting with trivial modifications.

[5] In [19], the largest constant used in specifying the endpoints of I is also considered. For our purpose here in this paper, the largest constant can be assumed to be ≤ 2 in all proofs.

tures that give rise to the differences in expressiveness. The necessity of a ‘new’ extension (such as the one in the next section) is justified by the fact that no known extension can lead to expressive completeness.

3.1 Definability of Time 0

Recall that $\mathsf{MTL}_{\mathsf{fut}}$ and $\mathsf{FO}[<,+1]$ have the same expressiveness over $[0,N)$-flows [16,18]. This result fails for the case of bounded timed words.

Proposition 1 (Corollary of [20, Section 8]). *MTL is strictly more expressive than $\mathsf{MTL}_{\mathsf{fut}}$ over $[0,N)$-timed words.*[6]

To account for this difference between the two semantics, observe that a distinctive feature of the continuous interpretation of $\mathsf{MTL}_{\mathsf{fut}}$ is exploited in [16,18]: 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 determine which unit interval in $[0,N)$ the current instant (where the relevant formula is being evaluated) is in. Unfortunately, since the duration of a given bounded timed word is not known *a priori*, this trick does not work for $\mathsf{MTL}_{\mathsf{fut}}$ in the pointwise semantics. For example, the formula $\Diamond_{=1}\mathbf{true}$ does not hold at any position in the $[0,2]$-timed word $\rho = (\sigma_0, 0)(\sigma_1, 0.5)$. However, the same effect can achieved in MTL by using past modalities. Let

$$\varphi_{i,i+1} = \overleftarrow{\Diamond}_{[i,i+1)}(\neg\overleftarrow{\Diamond}\mathbf{true})$$

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

Proposition 2. *$\mathsf{MTL}_{\mathsf{fut}}[\Phi_{unit}]$ is strictly more expressive than $\mathsf{MTL}_{\mathsf{fut}}$ over $[0,N)$-timed words.*

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

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

The models are illustrated in Figure 1, where each white box represents an event (at which no monadic predicate holds).

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

[6] The models constructed in [20, Section 8] are themselves bounded timed words.

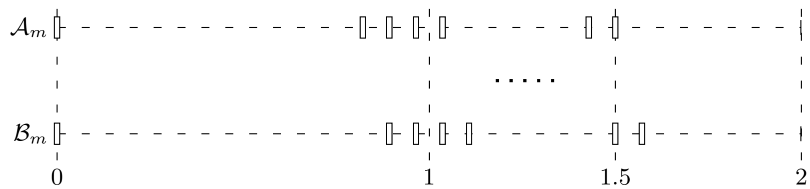

Fig. 1. Models $\mathcal{A}_m$ and $\mathcal{B}_m$

1). It follows from Theorem 1 that no $\mathsf{MTL_{fut}}$ formula of modal depth $\leq m$ distinguishes $\mathcal{A}_m$ and $\mathcal{B}_m$ while the formula

$$\Diamond_{(0,1)}\big(\varphi_{0,1} \wedge \bigcirc(\varphi_{0,1} \wedge \bigcirc\varphi_{0,1})\big)\,,$$

which says "there are three events with timestamps in $[0, 1)$", distinguishes $\mathcal{A}_m$ and $\mathcal{B}_m$ for any $m \in \mathbb{N}$ (when evaluated at position 0). $\square$

3.2 Past Modalities

The following proposition says that the conservative extension in the last subsection is not sufficient for obtaining expressive completeness: non-trivial nesting of future modalities and past modalities provides more expressiveness.

Proposition 3. *MTL is strictly more expressive than $\mathsf{MTL_{fut}}[\varPhi_{unit}]$ over $[0, N)$-timed words.*

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

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

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

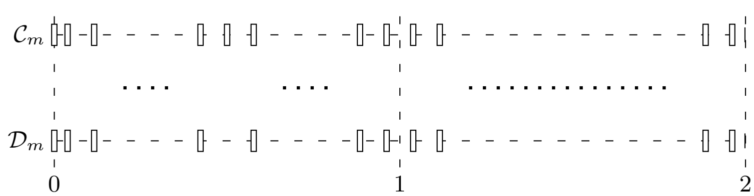

Fig. 2. Models $\mathcal{C}_m$ and $\mathcal{D}_m$

The models are illustrated in Figure 2, where each white box represents an event (at which no monadic predicate holds).

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. It follows from Theorem 1 that no $\mathsf{MTL}_{\mathsf{fut}}[\Phi_{unit}]$ formula of modal depth $\leq m$ distinguishes $\mathcal{C}_m$ and $\mathcal{D}_m$ while the formula

$$\Box_{(1,2)}(\diamondminus_{=1}\mathbf{true}),$$

which says "for each event with timestamp in $(1, 2)$, there is a corresponding event that is exactly 1 time unit earlier", distinguishes $\mathcal{C}_m$ and $\mathcal{D}_m$ for any $m \in \mathbb{N}$ (when evaluated at position 0). $\Box$

3.3 Counting Modalities

The modality $C_n(x, X)$ asserts that X holds at least at n points in the open interval $(x, x+1)$. The modalities C_n for $n \geq 2$ are called *counting modalities.* It is well-known that these modalities are inexpressible in MTL over $\mathbb{R}_{\geq 0}$-flows [10]. For this reason, they (or variants thereof) are often used to separate the expressiveness of various metric logics (cf., e.g., [2, 19, 20]). For example, the following property

- P holds at an event at time y in the future
- Q holds at an event at time $y' > y$
- R holds at an event at time $y'' > y' > y$
- Both the Q-event and the R-event are within $(1, 2)$ from the P-event

can be expressed as the $\mathsf{FO}[<, +1]$ formula

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

This formula has no equivalent in MTL over $\mathbb{R}_{\geq 0}$-timed words [19]. Indeed, it was shown recently that in the continuous semantics, MTL with counting modalities and their past counterparts (which we denote by $\mathsf{MTL}[\{C_n, \overleftarrow{C}_n\}_{n=2}^{\infty}]$) is expressively complete for $\mathsf{FO}[<, +1]$ [12]. However, counting modalities add no expressiveness to MTL in the time-bounded setting. To see this, observe that the following formula is equivalent to $\vartheta_{pqr}(x)$ over $[0, N)$-timed words (we make use of formulas in Φ_{unit} defined in Section 3.1):

$$\Diamond \Bigg(\bigvee_{i \in [0, N-1]} \Big(P \land \varphi_{i,i+1} \land \big(\Diamond_{>1} (Q \land \Diamond (R \land \varphi_{i+1,i+2})) \\ \lor \Diamond_{<2} (R \land \varphi_{i+2,i+3} \land \diamondminus (Q \land \varphi_{i+2,i+3})) \\ \lor (\Diamond_{>1} (Q \land \varphi_{i+1,i+2}) \land \Diamond_{<2} (R \land \varphi_{i+2,i+3})) \big) \Big) \Bigg).$$

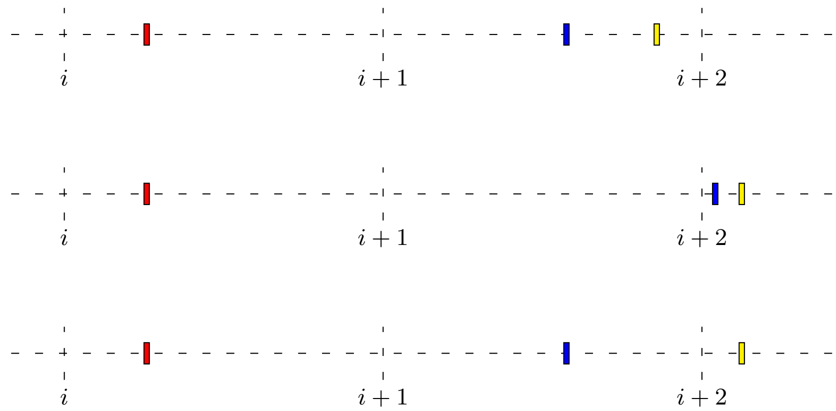

Fig. 3. Counting modalities is expressible

The three cases in the formula are illustrated in Figure 3 (from top to bottom) where times are measured relative to the current instant. Red, blue and yellow boxes represents P-events, Q-events and R-events respectively.

The same idea can readily be generalised to handle counting modalities and their past counterparts. We therefore have the following proposition.

Proposition 4. MTL *is expressively complete for* $\mathsf{MTL}[\{C_n, \overleftarrow{C}_n\}_{n=2}^{\infty}]$ *over* $[0, N)$*-timed words.*

3.4 Non-Local Properties: One Reference Point

Proposition 4 shows that part of the expressiveness hierarchy over $\mathbb{R}_{\geq 0}$-timed words collapses in the time-bounded setting. Nonetheless, MTL is still not expressive enough to capture all of $\mathsf{FO}[<, +1]$. Recall that another feature of the continuous interpretation of $\mathsf{MTL_{fut}}$ used in the proof in [16, 18] is that $\Diamond_{=k}\varphi$ holds at t *iff* φ holds at $t + k$. Suppose that we want to specify the following property over $\mathbf{P} = \{P, Q\}$ for some integer constant $a > 0$ (let the current instant be t_1):

- There is an event at time $t_2 > t_1 + a$ where Q holds
- P holds at all events in $(t_1 + a, t_2)$.

In the continuous semantics, the property can easily be expressed as the $\mathsf{MTL_{fut}}$ formula

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

over flows of the form f^ρ (over $\mathbf{P}' = P \cup \{P_\epsilon\}$. See Figure 4 for an example. Red boxes denote events at which P holds whereas blue boxes denote events at which Q holds. The formula φ_{cont1} holds at t_1 in the continuous semantics.

In essence, when the current time is t_1, the continuous interpretation of MTL allows one to speak of properties 'around' $t_1 + a$ regardless of whether there is

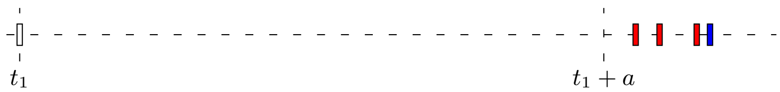

Fig. 4. φ_{cont1} holds at t_1 in the continuous semantics

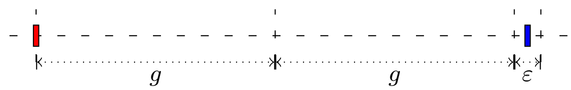

Fig. 5. A single segment in $\mathcal{E}_m$

an actual event at $t_1 + a$. The same is not readily possible with the pointwise interpretation of MTL if there is no event at $t_1 + a$. To handle this issue within the pointwise semantic framework, we introduce a relatively simple family of modalities $\mathcal{B}^{\rightarrow}_I$ ('Beginning') and their past versions $\mathcal{B}^{\leftarrow}_I$. They can be used to specify the *first* events in given intervals. To be precise, we define the modality that asserts "X holds at the first event in (a, b) relative to the current instant" as the following $\mathsf{FO}[<, +1]$ formula:

$$\mathcal{B}^{\rightarrow}_{(a,b)}(x, X) = \exists x' \Big(x < x' \wedge d(x, x') > a \wedge d(x, x') < b \wedge X(x') \\ \wedge \nexists x'' \big(x < x'' \wedge x'' < x' \wedge d(x, x'') > a \big) \Big) .$$

Now the property above can be defined as $\mathcal{B}^{\rightarrow}_{(a,\infty)}\big(Q \vee (P \,\mathcal{U}\, Q)\big)$ in the pointwise semantics. We refer to the extension of MTL with $\mathcal{B}^{\rightarrow}_I, \mathcal{B}^{\leftarrow}_I$ as $\mathsf{MTL}[\mathcal{B}^{\leftrightarrows}]$.

The following proposition states that this extension is indeed non-trivial.

Proposition 5. $\mathsf{MTL}[\mathcal{B}^{\leftrightarrows}]$ *is strictly more expressive than* MTL *over* $[0, N)$*-timed words.*

Proof. The proof we give here is inspired by a proof in [19, Section 5]. Given $m \in \mathbb{N}$, we describe models $\mathcal{E}_m$ and $\mathcal{F}_m$ that are indistinguishable by MTL formulas of modal depth $\leq m$ but distinguishing in $\mathsf{MTL}[\mathcal{B}^{\leftrightarrows}]$.

We first describe $\mathcal{E}_m$. Let $g = \frac{1}{2m+6}$ and pick $\varepsilon < \frac{g}{\frac{1}{g}-1}$. The first event (at time 0) satisfies $\neg P \wedge \neg Q$. Then, a sequence of overlapping segments (arranged as described below) starts at time $\frac{0.5}{2m+5}$. See Figure 5 for an illustration of a segment. Each segment consists of an event satisfying $P \wedge \neg Q$ (the red boxes) and an event satisfying $\neg P \wedge Q$ (the blue boxes). For ease of presentation we will refer to them as P-events and Q-events. If the P-event in the i^{th} segment is at time t, then its Q-event is at time $t + 2g + \frac{1}{2}\varepsilon$. All P-events in neighbouring segments are separated by $g - \frac{g}{\frac{1}{g}-1}$. We put a total of $4m + 12$ segments.

$\mathcal{F}_m$ is almost identical to $\mathcal{E}_m$ except the $(3m+9)^{th}$ segment. Let this segment start at t_{3m+9}. In $\mathcal{F}_m$, we move the corresponding Q-event to $t + 2g - \frac{1}{2}\varepsilon$ (see Figure 6). Note that there are P-events at time 0.5 in both models (in their $(m + 4)^{th}$ segment).

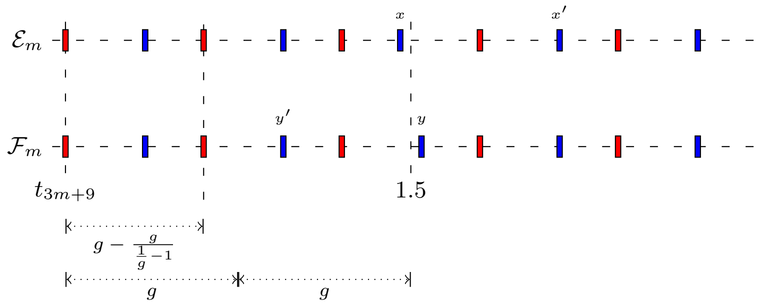

Fig. 6. Near the $(3m+9)^{th}$-segments in $\mathcal{E}_m$ and $\mathcal{F}_m$

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 6). It is easy to verify that no two events are separated by an integer distance. We say a configuration (i, j) is *identical* if $i = j$. For $i \geq 1$, 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 6. Duplicator *has a winning strategy for m-round* MTL *EF game on $\mathcal{E}_m, \mathcal{F}_m$ that starts from* $(0, 0)$. *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* $(3m+8-r) < seg(i_r), seg(j_r) < (3m+12+r)$.

We prove the proposition by induction on r. The idea is to try to make the resulting configurations identical. When this is not possible, *Duplicator* imitates what *Spoiler* does.

- *Base step.* The proposition holds trivially for $(i_0, j_0) = (0, 0)$.
- *Induction step.* Suppose that the claim holds for $r < m$. We prove it also holds for $r + 1$.
  - $(i_r, j_r) = (0, 0)$:
    Duplicator can always make (i_{r+1}, j_{r+1}) identical.
  - $(i_r, j_r) \neq (0, 0)$ 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 $\diamondsuit$-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 5 now follows from Proposition 6, Theorem 1, and the fact that $\mathcal{E}_m \models \diamondsuit(P \wedge \mathcal{B}^{\rightarrow}_{(1,2)} P)$ but $\mathcal{F}_m \not\models \diamondsuit(P \wedge \mathcal{B}^{\rightarrow}_{(1,2)} P)$. $\square$

3.5 Non-Local Properties: Two Reference Points

Adding modalities $\mathcal{B}^{\rightarrow}_I, \mathcal{B}^{\leftarrow}_I$ to MTL allows one to specify properties with respect to a distant time point even when there is no event at that point. However, the following proposition shows that this is still not enough for expressive completeness.

Proposition 7. $\mathsf{FO}[<, +1]$ *is strictly more expressive than* $\mathsf{MTL}[\mathcal{B}^{\leftrightarrows}]$ *over* $[0, N)$*-timed words.*

Proof. This is similar to a proof in [20, Section 7]. Given $m \in \mathbb{N}$, we construct two models as follows. Let

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

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

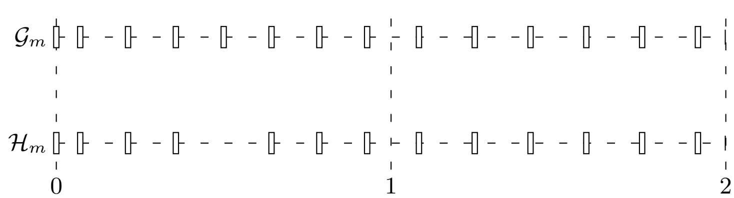

Fig. 7. Models $\mathcal{G}_m$ and $\mathcal{H}_m$ for $m = 2$

Figure 7 illustrates the models for the case $m = 2$ where white boxes represent events at which no monadic predicate holds. Observe that no two events are separated by an integer distance. We say that a configuration (i, j) is *synchronised* if they correspond to events with the same timestamp. Here we extend Theorem 1 with a $\mathcal{B}_I^{\rightarrow}$ move and its past counterpart. In this move, *Spoiler* picks (say) i'_r such that

– $\tau_{i'_r} - \tau_{i_r} \in I$ in ρ
– There is no position $i' < i'_r$ in ρ such that $\tau_{i'} - \tau_{i_r} \in I$,

and *Duplicator* must choose a position j'_r in ρ' such that j'_r is the first position in I relative to j_r in ρ'.

Proposition 8. Duplicator *has a winning strategy for m-round* $\mathsf{MTL}[\mathcal{B}^{\leftrightarrows}]$ *EF game on $\mathcal{G}_m, \mathcal{H}_m$ that starts from $(0, 0)$. In particular, she has a winning strategy such that for each round $0 \le r \le m$, 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* $(3m+5-r) < i_r, j_r < (3m+6+r)$.

We prove the proposition by induction on r. The idea, again, is to try to make the resulting configurations identical.

– *Base step.* The proposition holds trivially for $(i_0, j_0) = (0, 0)$.
– *Induction step.* Suppose that the claim holds for $r < m$. We prove it also holds for $r + 1$.
  - $(i_r, j_r) = (0, 0)$:
    Duplicator tries to make (i'_r, j'_r) synchronised. If *Spoiler* chooses $i'_r = m + 3$, *Duplicator* chooses $j'_r = m + 2$.
  - $(i_r, j_r) \neq (0, 0)$ 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 7 now follows from Proposition 8, our extended version of Theorem 1, and the fact that the $\mathsf{FO}[<, +1]$ formula

$$\begin{aligned}\exists x \Bigg(\nexists y\,(y < x) \wedge \exists x' \Bigg(& d(x, x') > 1 \wedge d(x, x') < 2 \\ & \wedge \exists x'' \Big(x' < x'' \wedge \nexists y'\,(x' < y' \wedge y' < x'') \\ & \qquad \wedge \nexists y''\,\big(d(x', y'') < 1 \wedge d(x'', y'') > 1\big)\Big)\Bigg)\Bigg)\end{aligned}$$

distinguishes $\mathcal{G}_m$ and $\mathcal{H}_m$ for any $m \in \mathbb{N}$. This formula asserts that there is a pair of neighbouring events in $(1, 2)$ such that there is no event between them if they are both mapped to exactly one time unit earlier. $\square$

One way to understand this phenomenon is to consider the arity of MTL operators. Let the current instant be t_1. Suppose that we want to specify the following property ($a > c > 0$):

- There is an event at $t_2 = t' + a > t_1 + a$ where Q holds
- P holds at all events in $\big(t_1 + c, t_1 + c + (t_2 - t_1 - a)\big)$.

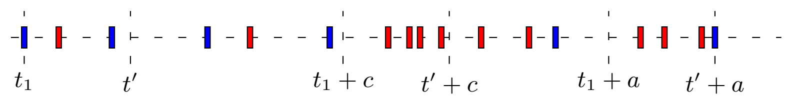

Fig. 8. φ_{cont2} holds at t_1 in the continuous semantics

See Figure 8 for an example where red boxes are P-events and blue boxes are Q-events. In the continuous semantics we can simply write

$$\varphi_{cont2} = \big(\Diamond_{=c}(P \vee P_\epsilon)\big)\, \mathcal{U}\, (\Diamond_{=a} Q)\,.$$

Observe how this formula (effectively) talks about properties around two points: $t_1{+}c$ and $t_1{+}a$. In the same vein, the following formula can be used to distinguish $\mathcal{G}_m$ and $\mathcal{H}_m$ (defined in the proof of Proposition 7) in the continuous semantics:

$$\varphi_{cont3} = \Diamond_{(1,2)}\big(\neg P_\epsilon \wedge (\overleftarrow{\Diamond}_{=1} P_\epsilon)\, \mathcal{U}\, (\neg P_\epsilon)\big)\,.$$

In the next section, we propose new modalities that add this ability to MTL in the pointwise semantics. We show later that this ability is exactly the missing piece of expressiveness.

4 New Modalities

4.1 Generalised 'Until' and 'Since'

We introduce a family of modalities which can be understood as generalisations of the usual 'Until' and 'Since' modalities. Let $I \subseteq (0, \infty)$ be an interval with endpoints in $\mathbb{N} \cup \{\infty\}$ and $c \in \mathbb{N}$. The formula $\varphi_1 \mathfrak{U}^c_I \varphi_2$ (using infix notation), when imposed at t_1, asserts that

- There is an event at t_2 where φ_2 holds and $t_2 - t_1 \in I$
- φ_1 holds at all events in $\Big(c, c + \Big(t_2 - \big(t_1 + \inf(I)\big)\Big)\Big)$ relative to t_1.

Formally, for $I = (a,b) \subseteq (0,\infty)$ and $a \geq c \geq 0$, we define the generalised 'Until' modality $\mathfrak{U}^c_{(a,b)}$ by the following $\mathsf{FO}[<,+1]$ formula:

$$\begin{aligned}\mathfrak{U}^c_{(a,b)}(x, X_1, X_2) = \exists x' \Big(x < x' \; \wedge d(x,x') > a \wedge d(x,x') < b \wedge X_2(x') \\ \wedge \forall x'' \left(x < x'' \wedge d(x,x'') > c \wedge x'' < x' \right. \\ \left. \wedge\, d(x',x'') > (a-c) \implies X_1(x'') \right) \Big) .\end{aligned}$$

Symmetrically, we define the generalised 'Since' modality $\mathfrak{S}^c_{(a,b)}$ for $I = (a,b)$ and $a \geq c \geq 0$:

$$\begin{aligned}\mathfrak{S}^c_{(a,b)}(x, X_1, X_2) = \exists x' \Big(x' < x \; \wedge d(x,x') > a \wedge d(x,x') < b \wedge X_2(x') \\ \wedge \forall x'' \left(x'' < x \wedge d(x,x'') > c \wedge x' < x'' \right. \\ \left. \wedge\, d(x',x'') > (a-c) \implies X_1(x'') \right) \Big) .\end{aligned}$$

We will refer to the logic obtained by adding these modalities to MTL as $\mathsf{MTL}[\mathfrak{U},\mathfrak{S}]$. Note that the usual 'Until' and 'Since' modalities can be written in terms of generalised modalities. For instance,

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

4.2 More Liberal Bounds

In the definition of modalities $\mathfrak{U}^c_I$ and $\mathfrak{S}^c_I$ in the last subsection, we stressed that $I \subseteq (0,\infty)$ and $\inf(I) \geq c \geq 0$. This is because more liberal usage of bounds are indeed merely syntactic sugar. For instance, one may define a modality $\mathfrak{U}^{10}_{(2,5)}(x, X_1, X_2)$ to assert the following property when imposed at t_1:

- There is an event at t_2 where φ_2 holds and $t_2 - t_1 \in (2,5)$
- φ_1 holds at all events in $(10,13)$ relative to t_1.

This can be expressed in $\mathsf{FO}[<,+1]$ as the formula

$$\begin{aligned}\exists x' \Big(x < x' \; \wedge d(x,x') > 2 \wedge d(x,x') < 5 \wedge X_2(x') \\ \wedge \forall x'' \left(x < x'' \wedge d(x,x'') > 10 \wedge d(x',x'') < 8 \implies X_1(x'') \right) \Big) .\end{aligned}$$

However, this definition is not necessary as the formula above is indeed equivalent to

$$\Diamond_{(2,5)}\varphi_2 \wedge \neg\Big((\neg\varphi_2)\,\mathfrak{U}^2_{(10,13)}\,\left(\neg\varphi_1 \wedge \neg(\diamondminus_{=8}\varphi_2)\right)\Big)$$

if we substitute φ_1, φ_2 for X_1, X_2. This can be further generalised as follows.

Proposition 9. *The property (imposed at t_1):*

- *There is an event at t_2 where φ_2 holds and $t_2 - t_1 \in (a,b)$*

– φ_1 *holds at all events in* $\Big(c, c + \big(t_2 - (t_1 + a)\big)\Big)$ *relative to* t_1

with $c \in \mathbb{Z}$*,* $I = (a, b) \subseteq (-\infty, \infty)$ *and* $a \in \mathbb{Z}$ *can be expressed in modalities defined in Section 4.1.*

Proof. (*By example*) For example, the formula $\varphi_1 \, \mathfrak{U}^{-7}_{(5,10)} \, \varphi_2$ (with the intended meaning as described above) is equivalent to

$$\neg\varphi_2 \mathfrak{U}^{5}_{(5,10)} \big(\varphi_2 \wedge (\varphi_1 \mathfrak{S}^{12}_{(5,10)} \mathbf{true})\big) \wedge \Big((\mathbf{false}\, \mathfrak{U}^{0}_{(5,10)} \varphi_2) \vee \Big(\varphi' \mathcal{U} \big((\mathbf{false}\, \mathfrak{U}^{0}_{(5,10)} \varphi_2) \wedge \varphi'\big)\Big)\Big)$$

where $\varphi' = \varphi_1 \, \mathfrak{S}^{7}_{(0,5)} \, \big(\mathbf{true} \wedge \neg(\diamondsuit_{=7} \neg\varphi_1)\big)$. This allows us to write, e.g.,

$$\varphi_1 \, \mathfrak{U}^{5}_{(-7,-2)} \, \varphi_2 = \diamondsuit_{(2,7)} \varphi_2 \wedge \neg\Big((\neg\varphi_2) \, \mathfrak{U}^{-7}_{(5,10)} \, \big(\neg\varphi_1 \wedge \neg(\diamondsuit_{=12} \varphi_2)\big)\Big) \, .$$

As another example, $\varphi_1 \, \mathfrak{S}^{-15}_{(-10,-5)} \, \varphi_2$ is equivalent to the disjunction of

$$\mathbf{false} \, \mathfrak{U}^{0}_{(-10,-5)} \, \big(\varphi_2 \wedge (\varphi_1 \, \mathfrak{S}^{5}_{(-10,-5)} \, \mathbf{true})\big)$$

and

$$\diamondsuit_{(5,10)} \varphi_2 \wedge \Bigg(\varphi'' \, \mathcal{U} \, \bigg(\Big(\mathbf{false} \, \mathfrak{U}^{0}_{(-10,-5)} \, \big(\varphi_2 \wedge (\varphi_1 \, \mathfrak{S}^{5}_{(-10,-5)} \, \mathbf{true})\big)\Big) \wedge \varphi'' \bigg)\Bigg)$$

where $\varphi'' = \varphi_1 \, \mathfrak{S}^{15}_{(0,5)} \, \big(\mathbf{true} \wedge \neg(\diamondsuit_{=15} \neg\varphi_1)\big)$. The remaining cases can be handled with similar ideas. ☐

We can now give an $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ formula that distinguishes, in the pointwise semantics, the models $\mathcal{G}_m$ and $\mathcal{H}_m$ in Section 3.5:

$$\diamondsuit_{(1,2)} \big(\mathbf{true} \wedge (\mathbf{false} \, \mathfrak{U}^{-1}_{(0,\infty)} \, \mathbf{true})\big) \, .$$

This formula is, in essence, very similar to the formula φ_{cont3} defined in Section 3.5, which distinguishes $\mathcal{G}_m$ and $\mathcal{H}_m$ in the continuous semantics.

5 The Translation

We give a translation from an arbitrary $\mathsf{FO}[<, +1]$ formula with one free variable into an equivalent $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ formula (over $[0, N)$-timed words). Our proof strategy is similar to that in [16]: we eliminate the metric by introducing extra predicates, convert to LTL, and then replace the new predicates by their equivalent $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ 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 [16] and, additionally, $\overline{\mathbf{Q}} = \{Q_i \mid 0 \leq i \leq N-1\}$. Intuitively, $P_i(x)$ holds (for $x \in [0,1)$) iff $P \in \mathbf{P}$ holds at time $i+x$ in the corresponding $[0,N)$-timed word, and $Q_i(x)$ holds iff there is an event at time $i+x$ in the corresponding $[0,N)$-timed word, regardless of whether any $P \in \mathbf{P}$ holds there. Let $\vartheta_{event} = \forall x \left(\bigvee_{i\in[0,N-1]} Q_i(x)\right) \wedge \forall x \left(\bigwedge_{i\in[0,N-1]} \left(P_i(x) \implies Q_i(x)\right)\right)$ and $\vartheta_{init} = \exists x \left(\nexists x' \, (x' < x) \wedge Q_0(x)\right)$. There is an obvious 'stacking' bijection (indicated by overlining) between $[0,N)$-timed words over $\Sigma_{\mathbf{P}}$ and $[0,1)$-timed words over $\Sigma_{\overline{\mathbf{P}}\cup\overline{\mathbf{Q}}}$ satisfying $\vartheta_{event} \wedge \vartheta_{init}$. For example, the stacked counterpart of the $[0,2)$-timed word

$$(\{A\},0)(\{A,C\},0.3)(\{B\},1)(\{B,C\},1.5)$$

with $\mathbf{P} = \{A,B,C\}$ is a $[0,1)$-timed word:

$$(\{Q_0,Q_1,A_0,B_1\},0)(\{Q_0,A_0,C_0\},0.3)(\{Q_1,B_1,C_1\},0.5)\,.$$

Let $\vartheta(x)$ be an $\mathsf{FO}[<,+1]$ formula with one free variable and in which each quantifier uses a fresh new variable. Without loss of generality, we assume that $\vartheta(x)$ contains only existential quantifiers (this can be achieved by syntactic rewriting). Replace the formula by

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

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

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

Note that we do not actually have the $+1$ function in our structures; it only serves as annotation here and will be removed later, e.g., $x'+k$ means that $Q_k(x')$ holds. We 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 $k_1 < k_2$
  - $\neg$**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 $k_1 - k_2 = 2$
  - $\neg$**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 $\Sigma_{\mathbf{P}}$ and its stacked counterpart $\overline{\rho}$, the following holds:

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

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

By Kamp's theorem [14], $\overline{\vartheta}(x)$ is equivalent to an $\mathsf{LTL}[\mathcal{U}, \mathcal{S}]$ formula $\overline{\varphi}$ of the following form:

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

5.2 From Non-Metric to Metric

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

Proposition 10. *Let $\overline{\psi}$ be a subformula of $\overline{\varphi}_i$ for some $i \in [0, N-1]$. There is an $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ formula ψ such that for any $[0, N)$-timed word ρ, $t \in \rho$ and $frac(t) = \overline{t} \in \overline{\rho}$, we have*

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

Proof. The $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ formula ψ is constructed inductively as follows:

- *Base step.* Consider the following cases:
  - $\overline{\psi} = P_j$: Let
  $$\psi = (\varphi_{0,1} \wedge \Diamond_{=j} P) \vee \ldots \vee (\varphi_{j,j+1} \wedge P) \vee \ldots \vee (\varphi_{N-1,N} \wedge \overleftarrow{\Diamond}_{=((N-1)-j)} P) .$$
  - $\overline{\psi} = Q_j$: Similarly we let
  $$\psi = (\varphi_{0,1} \wedge \Diamond_{=j} \mathbf{true}) \vee \ldots \vee (\varphi_{j,j+1} \wedge \mathbf{true}) \vee \ldots \vee (\varphi_{N-1,N} \wedge \overleftarrow{\Diamond}_{=((N-1)-j)} \mathbf{true}) .$$
- *Induction step.* The case for boolean operations is trivial and hence omitted.
  - $\overline{\psi} = \overline{\psi}_1 \, \mathcal{U} \, \overline{\psi}_2$: By IH we have ψ_1 and ψ_2. Let
  $$\psi^{j,k,l} = \psi_1 \, \mathfrak{U}^k_{(j,j+1)} \, (\psi_2 \wedge \varphi_{l,l+1}) .$$

  The desired formula is

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

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

The claim holds by a straightforward induction on the structure of $\overline{\psi}$ and ψ. $\square$

Construct corresponding formulas φ_i for each $\overline{\varphi}_i$ using the proposition above. Substitute them into $\overline{\varphi}$ 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 11. *For any* $[0, N)$*-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 10. $\square$

We are now ready to state our main result.

Theorem 2. $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ *is expressively complete for* $\mathsf{FO}[<, +1]$ *over* $[0, N)$*-timed words.*

6 Time-Bounded Verification

We now show that the timed-bounded satisfiability and model-checking problems for $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ are EXPSPACE-complete.

Theorem 3. *The time-bounded satisfiability problem for* $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ *is* EXPSPACE-*complete.*

Proof. Recall from Section 5.1 that an (untimed) finite word that satisfies ϑ_{init} and ϑ_{event} can be embedded into a $[0, 1)$-timed word $\overline{\rho}$. Given an $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ formula φ, we can construct an equisatisfiable LTL formula φ' of exponential size. The idea is that for each subformula ψ of φ and every $i \in [0, N)$, we introduce a monadic predicate F_i^ψ. We then add suitable subformulas in φ' to ensure that F_i^ψ holds at time $\overline{t} \in \overline{\rho}$ iff ψ holds at time $t \in \rho$ (where ρ is the unstacked counterpart of $\overline{\rho}$ and $frac(t) = \overline{t}$).

The construction in [16] carries over to our pointwise case except for the new modalities which can be handled as follows. For example, for $i \leq N - 4$ and a subformula $P \,\mathfrak{U}^1_{2,3}\, Q$ of φ, we add to φ' the following subformula as a conjunct:

$$\begin{aligned} F_i^{P \mathfrak{U}^1_{(2,3)} Q} \leftrightarrow {} & \big((Q_{i+1} \to F^P_{i+1}) \,\mathcal{U}\, F^Q_{i+2}\big) \\ & \vee \Big(\Box(Q_{i+1} \to F^P_{i+1}) \wedge \diamondminus\big(F^Q_{i+3} \wedge \boxminus(Q_{i+2} \to F^P_{i+2})\big)\Big)\,. \end{aligned}$$

EXPSPACE-hardness follows from the proof of EXPSPACE-hardness of $\mathsf{Bounded\text{-}MTL}$ in [3] as it readily carries over to our bounded-time setting. $\square$

Since the time-bounded model-checking problem and satisfiability problem are interreducible [9, 16], we have the following theorem.

Theorem 4. *The time-bounded satisfiability problem for timed automata against* $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$ *is* EXPSPACE-*complete.*

7 Conclusion

Our main result is that over bounded timed words, MTL extended with our new modalities 'generalised until' and 'generalised since' is expressively complete for $\mathsf{FO}[<, +1]$. Along the way we obtain a strict hierarchy of metric temporal logics, based on their expressiveness over bounded timed words:

$$\mathsf{MTL}_{\mathsf{fut}} \subsetneq \mathsf{MTL}_{\mathsf{fut}}[\Phi_{unit}] \subsetneq \mathsf{MTL} \subsetneq \mathsf{MTL}[\mathcal{B}^{\leftrightarrows}] \subsetneq \mathsf{MTL}[\mathfrak{U}, \mathfrak{S}] = \mathsf{FO}[<, +1].$$

The proposed modalities $\mathfrak{U}$ and $\mathfrak{S}$ are not very intuitive. However, as we proved that a simpler variant of these modalities are strictly less expressive over bounded timed words, it is unlikely that a reasonable expressively complete extension of MTL, if exists, would be "simpler" than our extensions.

A possible future direction would be to investigate whether a similar expressive completeness result can be obtained in the case of $\mathbb{R}_{\geq 0}$-timed words. This is likely to require a separation theorem (in the style of [13]) that works in the pointwise semantics.

The idea of considering bounded time domains is useful in dealing with other formalisms as well. For example, new decidability and complexity results have been obtained for hybrid automata in time-bounded settings [4, 5]. It would be interesting to investigate the whether our results can be combined with such new developments in practical applications. For example, an implication of our result to the model-checking and satisfiability problems (of timed automata against $\mathsf{MTL}[\mathfrak{U}, \mathfrak{S}]$) is that traditional, highly-matured LTL-based verification tools such as SPIN [11] and SPOT [7] can easily be utilised in implementations.

References

1. Alur, R., Dill, D.: A theory of timed automata. Theoretical Computer Science 126(2), 183–235 (1994)
2. 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)
3. Bouyer, P., Markey, N., Ouaknine, J., Worrell, J.: The cost of punctuality. In: LICS. pp. 109–120. IEEE Computer Society (2007), `http://doi.ieeecomputersociety.org/10.1109/LICS.2007.49`
4. Brihaye, T., 0001, L.D., Geeraerts, G., Ouaknine, J., Raskin, J.F., Worrell, J.: On reachability for hybrid automata over bounded time. In: Aceto, L., Henzinger, M., Sgall, J. (eds.) ICALP (2). Lecture Notes in Computer Science, vol. 6756, pp. 416–427. Springer (2011), `http://dx.doi.org/10.1007/978-3-642-22012-8`
5. Brihaye, T., Doyen, L., Geeraerts, G., Ouaknine, J., Raskin, J.F., Worrell, J.: Time-bounded reachability for monotonic hybrid automata: Complexity and fixed points. In: Dang-Van, H., Ogawa, M. (eds.) Proceedings of the 11th International Symposium on Automated Technology for Verification and Analysis (ATVA'13). Lecture Notes in Computer Science, vol. 8172, pp. 55–70. Springer (Oct 2013), `http://www.lsv.ens-cachan.fr/Publis/PAPERS/PDF/BDGORW-atva13.pdf`
6. D'Souza, D., Holla, R., Vankadaru, D.: On the expressiveness of TPTL in the pointwise and continuous semantics. Unpublished manuscript (2007)

7. Duret-Lutz, A., Poitrenaud, D.: SPOT: An extensible model checking library using transition-based generalized Büchi automata. In: Proceedings of MASCOTS 2004. pp. 76–83. IEEE Computer Society Press (2004)
8. Etessami, K., Wilke, T.: An until hierarchy for temporal logic. In: LICS. pp. 108–117. IEEE Computer Society (1996), http://doi.ieeecomputersociety.org/10.1109/LICS.1996.561310
9. Henzinger, T.A., Raskin, J.F., Schobbens, P.Y.: The regular real-time languages. In: Larsen, K.G., Skyum, S., Winskel, G. (eds.) ICALP. Lecture Notes in Computer Science, vol. 1443, pp. 580–591. Springer (1998), http://dx.doi.org/10.1007/BFb0055086
10. Hirshfeld, Y., Rabinovich, A.: Expressiveness of metric modalities for continuous time. Logical Methods in Computer Science 3(1) (2007)
11. Holzmann, G.J.: The model checker SPIN. IEEE Transactions on Software Engineering 23(5), 279–295 (1997)
12. 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)
13. 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)
14. Kamp, J.: Tense logic and the theory of linear order. Ph.D. thesis, University of California, Los Angeles (1968)
15. Koymans, R.: Specifying real-time properties with metric temporal logic. Real-Time Systems 2(4), 255–299 (1990)
16. Ouaknine, J., Rabinovich, A., Worrell, J.: Time-bounded verification. In: Proceedings of CONCUR 2009. LNCS, vol. 5710, pp. 496–510. Springer (2009)
17. Ouaknine, J., Worrell, J.: Some recent results in metric temporal logic. In: Proceedings of FORMATS 2008. LNCS, vol. 5215, pp. 1–13. Springer (2008)
18. Ouaknine, J., Worrell, J.: Towards a theory of time-bounded verification. In: Proceedings of ICALP 2010. LNCS, vol. 6199, pp. 22–37. Springer (2010)
19. 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)
20. 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)
21. Wilke, T.: Specifying timed state sequences in powerful decidable logics and timed automata. In: Proceedings of FTRTFT 1994. LNCS, vol. 863, pp. 694–715. Springer (1994)