MTL with Bounded Variability: Decidability and Complexity

Carlo A. Furia and Matteo Rossi

May 2008

Abstract

This paper investigates the properties of Metric Temporal Logic (MTL) over models in which time is dense but phenomena are constrained to have *bounded variability*. Contrary to the case of generic dense-time behaviors, MTL is proved to be fully decidable over models with bounded variability, if the variability bound is given. In these decidable cases, MTL complexity is shown to match that of simpler decidable logics such as MITL. On the contrary, MTL is undecidable if all behaviors with variability bounded by some generic constant are considered, but with an undecidability degree that is lower than in the case of generic behaviors.

Contents

1 Introduction **3**
1.1 Related Work 4

2 Words and Behaviors: A Semantic Zoo **5**
2.1 Words and Behaviors 5
2.2 Finite Variability and non-Zenoness 6
2.3 Bounded Variability and Non-Berkeleyness 7
2.3.1 Bounded Variability 7
2.3.2 Non-Berkeleyness 7
2.3.3 Relations Among Classes 7

3 MTL and Its Relatives **10**
3.1 MITL and MTL 10
3.2 QTL and QITL: Decidable Extensions of MITL 12

4 Syntactic Definition of Regularity Constraints **14**
4.1 From Non-Berkeleyness to Bounded Variability 14
4.2 Describing Sequences of Transitions 16
4.3 Syntactic Characterizations 17
4.3.1 Non-Berkeleyness 18
4.3.2 Bounded Variability 19

5 Decidability Results **19**
5.1 MTL over Non-Berkeley Behaviors 20
5.1.1 From Flat MTL to MITL 20
5.1.2 From MTL to QITL 20
5.1.3 Decidability of MTL over Non-Berkeley Behaviors 21
5.2 MTL over Bounded Variably Behaviors 22
5.3 MTL over Words 22

6 Related Results **23**
6.1 Expressiveness of MTL over non-Berkeley 23
6.2 Complexity of MTL over Non-Berkeley 25

7 Undecidability Results **27**

8 Summary **29**

1 Introduction

The designer of formal notations faces a perennial trade-off between expressiveness and complexity: on the one hand notations with high expressive power allow users to formalize complex behaviors with naturalness; on the other hand expressiveness usually comes with a significant price to pay in terms of complexity of the verification problem. This results in a continual search for the "best" compromise between these diverging features.

A paradigmatic instance of this general problem is the case of real-time temporal logics. Experience with real-time concurrent systems suggests that dense (or continuous) sets are a natural and effective modeling choice for the time domain. Also, Metric Temporal Logic (MTL) is often regarded as a suitable and natural extension of "classical" Temporal Logic to deal with real-time requirements. However, MTL is well-known to be undecidable over dense time domains [AH93].[1] In the literature, two main compromises have been adopted to overcome this impasse. One consists in the semantic accommodation of adopting the coarser discrete — rather than dense — time [AH93]. The other adopts the syntactic concession of restricting MTL formulas to a subset known as MITL [AFH96]. More recently other syntactic adjustments have been studied [BMOW07].

In this paper we investigate other semantic compromises, in particular the use of models where time is dense but events are constrained to have only a *bounded variability*, i.e., their frequency of occurrence over time is bounded by some finite constant. We show that MTL is fully decidable over such behaviors when the maximum variability rate is fixed *a priori*; in such cases we are also able to show that the complexity of decidability is the same as for the less expressive logic MITL, i.e., complete for **EXPSPACE**. On the contrary, if all behaviors with bounded variability are considered together, MTL becomes undecidable, but with a "lesser degree" of undecidability compared to the case of unconstrained behaviors. Our decidability results are based on the possibility of expressing certain features of bounded variability in the expressive decidable temporal logics of [HR04]. Although the focus of this paper is on the more expressive *behavior* semantic model (also called signal, timed interval sequence, or trajectory) which is more expressive [DP07] but requires more sophisticated techniques, one can show that the same decidability and complexity results hold in the timed word case as well (where they were already partly implied by the results in [Wil94]).

Paper outline. The rest of the paper is organized as follows. The next subsection summarizes the related works that are most closely connected to the results in the present paper. Section 2 introduces the various semantic models considered in this paper, namely behaviors (and words) with bounded variability, as well as those with a similar constraint called non-Berkeleyness. Section 3 introduces the temporal logics MTL, MITL, and a decidable extension of the latter called QITL. Section 4 shows how the semantic classes introduced beforehand can be syntactically characterized through some of the logics of Section 3. Section 5 proves the decidability results of the paper, Section 6 discusses the complexity of the decidable logics, whereas Section 7 proves some undecidability

[1] With a few partial exceptions that will be discussed in the following.

results. Finally, Section 8 summarizes the results of the paper.

1.1 Related Work

The complexity, decidability, and expressiveness of MTL over standard discrete and dense time models are well-known since the seminal work of Alur and Henzinger [AH93] (which popularized the propositional subset of Koyman's original notation [Koy90]). In [AH93] MTL is shown to be decidable over discrete time, with an **EXPSPACE**-complete decidability problem, and undecidable over dense time, with a Σ_1^1-complete decidability problem. These results hold regardless of whether a timed word or timed signal time model is assumed, with a peculiar, but significant exception: in a recent, unexpected, result, Ouaknine and Worrell showed that MTL is fully decidable over *finite* dense-timed words, if only future modalities are considered [OW07]. The practical usefulness of this result is unfortunately plagued by the prohibitively high nonprimitive-recursive complexity of the corresponding decidability problem.

In another very influential paper [AFH96], Alur, Feder, and Henzinger showed that disallowing the expression of punctual (i.e., exact) time distances in MTL formulas renders the language fully decidable over dense time models. The corresponding MTL subset is called MITL and has an **EXPSPACE**-complete decidability problem. The decision procedure in [AFH96] is based on a complex translation into timed automata; similar, but simpler, automata-based techniques have been studied by Maler et al. [MNP06].

Hirshfeld and Rabinovich have reconsidered the work on MITL from a broader, more foundational, perspective built upon the standard timed behavior model [HR04]. Besides providing simpler decision procedures and proofs for a real-time temporal logic with the same expressive power as MITL, they have probed to what extent MITL can be made more expressive without giving up decidability. This lead to the introduction of the very expressive, yet decidable, monadic logic Q2MLO, and of the corresponding TLC temporal logic. In a nice analogy with classical results on linear temporal logic [GHR94], TLC is expressively complete for all of Q2MLO (hence it subsumes MITL), and it has a **PSPACE**-complete decidability problem (or **EXPSPACE**-complete assuming a succinct encoding of constants used in formulas as it is customary in the majority of MITL literature).

It is clear that TLC and MTL have incomparable expressive power; in particular the former disallows the expression of exact time distances. However, Bouyer et al. [BMOW07] have shown that it is possible to devise significantly expressive MTL fragments that are fully decidable (with **EXPSPACE** complexity) even if punctuality requirements are allowed to some extent. Also, for brevity we omit the summary of other, related complexity results for decidable real-time temporal logics over dense time domains recently developed by Lutz et al. [LWW07].

Dense-timed words where the maximum number of events in a unit interval is fixed have been introduced by Wilke in [Wil94]. More precisely, timed words over Σ where there are at most k positions over any unit interval are denoted by $\mathrm{TSS}_k(\Sigma)$ and called words of *bounded variability* k; in the following we introduce the class $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{k,1}$ that can be seen to correspond to $\mathrm{TSS}_k(\Sigma)$. Wilke showed that, for every k, the monadic logic of distances $\mathcal{L}\mathrm{d}(\Sigma)$ is fully decidable over $\mathrm{TSS}_k(\Sigma)$. Wilke's results are based on translation into the monadic fragment

$\overleftrightarrow{\mathcal{L}\mathrm{d}}(\Sigma)$, which ultimately corresponds to timed automata; also, they subsume the decidability of MTL over the same models. In this paper, we extend and generalize Wilke's result, and we discuss the complexity of the corresponding models.

The corresponding notion of dense-time *behaviors* with bounded variability has been introduced by Fränzle in [Frä96] (where they are called *trajectories of n-bounded variability*). Fränzle shows that full Duration Calculus is undecidable even over such restricted behaviors, while some syntactic subsets of it become decidable; the decidability proofs exploit a characterization of certain behaviors with bounded variability by means of timed regular expressions [ACM02].

In previous work [FR06, FPR08], we introduced the notion of non-Berkeleyness: a dense-time behavior is non-Berkeley for some $\delta > 0$ if δ time units elapse between any two consecutive state transitions. In this paper we show that this notion is similar, but different, than the notion of bounded variability; we also introduce a corresponding definition of non-Berkeleyness for timed words.

2 Words and Behaviors: A Semantic Zoo

The symbols $\mathbb{Z}$, $\mathbb{Q}$, and $\mathbb{R}$ denote the sets of integer, rational, and real numbers, respectively. For a set $\mathbb{S}$, $\mathbb{S}_{\sim c}$ with $\sim$ one of $<, \leq, >, \geq$ and $c \in \mathbb{S}$ denotes the subset $\{s \in \mathbb{S} \mid s \sim c\} \subseteq \mathbb{S}$; for instance $\mathbb{Z}_{\geq 0} = \mathbb{N}$ denotes the set of nonnegative integers (i.e., naturals).

An *interval* I of a set $\mathbb{S}$ is a convex subset $\langle l, u \rangle$ of $\mathbb{S}$ with $l, u \in \mathbb{S}$, $\langle$ one of $($, $[$, and $\rangle$ one of $)$, $]$. An interval is *empty* iff it contains no points; an interval is *punctual* (or singular) iff $l = u$ and the interval is closed (i.e., it contains exactly one point). The *length* of an interval is given by $|I| = \max(u - l, 0)$. $-I$ denotes the interval $\langle -u, -l \rangle$, and $I \oplus t = t \oplus I$ denotes the interval $\langle t + l, t + u \rangle$, for any $t \in \mathbb{S}$.

For a finite (ordered) sequence $\mathbf{x} = x_1, x_2, \ldots, x_n$, let $[\mathbf{x}]$ denote the sequence obtained by removing (consecutive) duplicate elements from $\mathbf{x}$. Correspondingly, we define the *length* $|\mathbf{x}|$ of $\mathbf{x}$ as $|\mathbf{x}| = n$ and its *cardinality* $\langle \mathbf{x} \rangle$ as the length $\langle \mathbf{x} \rangle = |[\mathbf{x}]|$ of $[\mathbf{x}]$. The intersection between a sequence and a set is a sequence obtained by projecting all symbols not in the set out of the sequence.

2.1 Words and Behaviors

The two most popular models of real-time behavior [AH93, ACM02] are the *timed word* (also called timed state sequence [Wil94]) and the *timed behavior* (also called Boolean signal [MNP05], timed interval sequence [AFH96], or trajectory [Frä96]).

Let $\mathbb{T}$ be a time domain; in this paper we are interested in dense time domains, and in particular $\mathbb{R}$ and its mono-infinite subset $\mathbb{R}_{\geq 0}$. Also, let Σ be a set of atomic propositions.

Behaviors. A *(timed) behavior* over timed domain $\mathbb{T}$ and alphabet Σ is a function $b : \mathbb{T} \to 2^{\Sigma}$ which maps every time instant $t \in \mathbb{T}$ to the set of propositions $b(t) \in 2^{\Sigma}$ that hold at t. The set of all behaviors over time domain $\mathbb{T}$ and alphabet Σ is denoted by $\overline{\mathcal{B}\Sigma\mathbb{T}}$.

For a behavior b let $\tau(b)$ denote the ordered (multi)set of its discontinuity points, i.e., $\tau(b) = \{x \in \mathbb{T} \mid b(x) \neq \lim_{t \to x^-} b(t) \vee b(x) \neq \lim_{t \to x^+} b(t)\}$, where each point that is both a right- and a left-discontinuity appears twice in $\tau(b)$. If $\tau(b)$ is discrete, we can represent it as an ordered sequence (possibly unbounded to $\pm\infty$); it will be clear from the context whether we are treating $\tau(b)$ as a sequence or as a set. Elements in $\tau(b)$ are called the *change* (or *transition*) instants of b. $\tau(b)$ can be unbounded to $\pm\infty$ only if $\mathbb{T}$ has the same property.

Words. An *infinite (timed) word* over time domain $\mathbb{T}$ and alphabet Σ is a sequence $(\Sigma \times \mathbb{T})^\omega \ni (\boldsymbol{\sigma}, \mathbf{t}) = (\sigma_0, t_0)(\sigma_1, t_1) \cdots$ such that: (1) for all $k \in \mathbb{N} : \sigma_k \in 2^\Sigma$, and (2) the sequence $\mathbf{t}$ of timestamps is strictly monotonically increasing. Every element (σ_n, t_n) in a word denotes that the propositions in the set σ_n hold at time t_n. The set of all infinite timed words over time domain $\mathbb{T}$ and alphabet Σ is denoted by $\overline{\mathcal{T}\Sigma\mathbb{T}^\omega}$. *Finite timed words* over time domain $\mathbb{T}$ and alphabet Σ are defined similarly as finite sequences in $(\Sigma \times \mathbb{T})^*$ and collectively denoted by $\mathcal{T}\Sigma\mathbb{T}^*$. Also, the set of all finite timed words of *length up to n* is denoted by $\mathcal{T}\Sigma\mathbb{T}^n = \{(\boldsymbol{\sigma}, \mathbf{t}) \in \mathcal{T}\Sigma\mathbb{T}^* \mid |\mathbf{t}| \leq n\}$.

2.2 Finite Variability and non-Zenoness

Since one is typically interested only in behaviors that represent physically meaningful behaviors, it is common to assume some regularity requirements on words and behaviors. In particular, it is customary to assume *non-Zenoness*, also called *finite variability* [HR04].

Non-Zeno behaviors. A behavior $b \in \overline{\mathcal{B}\Sigma\mathbb{T}}$ is non-Zeno iff $\tau(b)$ has no accumulation points; non-Zeno behaviors are denoted by $\mathcal{B}\Sigma\mathbb{T}$.

It should be clear that every non-Zeno behavior can be represented through a canonical countable sequence of adjacent intervals of $\mathbb{T}$ such that b is constant on every such interval. Namely, for $b \in \mathcal{B}\Sigma\mathbb{T}$, $\iota(b)$ is an ordered sequence of intervals $\iota(b) = \{I_i = \langle^i l_i, u_i \rangle^i \mid i \in \mathbb{I}\}$ such that:

1. *(cardinality of $\iota(b)$)* $\mathbb{I}$ is an interval of $\mathbb{Z}$ with cardinality $|\tau(b)| + 1$ (in particular, $\mathbb{I}$ is finite iff $\tau(b)$ is finite, otherwise $\mathbb{I}$ is denumerable);
2. *(partitioning of $\mathbb{T}$)* the intervals in $\iota(b)$ form a partition of $\mathbb{T}$;
3. *(intervals change at transition points)* for all $i \in \mathbb{I}$ we have $\tau_i = u_i = l_{i+1}$;
4. *(b constant over intervals)* for all $i \in \mathbb{I}$, for all $t_1, t_2 \in I_i$ we have $b(t_1) = b(t_2)$.

Note that $\iota(b)$ is unique for any fixed $\tau(b)$ or, in other words, is unique up to translations of interval indices. Transitions at instants τ_i corresponding to singular intervals I_i are called *pointwise* (or *punctual*) transitions.

Non-Zeno words. An *infinite* word $w \in \overline{\mathcal{T}\Sigma\mathbb{T}^\omega}$ is non-Zeno iff the sequence $\mathbf{t}$ of timestamps is diverging; non-Zeno infinite timed words are denoted by $\mathcal{T}\Sigma\mathbb{T}^\omega$. On the other hand, every *finite* timed word is non-Zeno by definition.

2.3 Bounded Variability and Non-Berkeleyness

In this paper we investigate behavior and words subject to regularity requirements that are stricter than non-Zenoness. In this section we introduce the two closely related — albeit different — notions of *bounded variability* and *non-Berkeleyness*.

2.3.1 Bounded Variability

A *behavior* $b \in \mathcal{B}\Sigma\mathbb{T}$ has *variability bounded* by k, δ for $k \in \mathbb{N}_{>0}, \delta \in \mathbb{R}_{>0}$ iff it has at most k transition points over every open interval of size δ. The set of all behaviors in $\mathcal{B}\Sigma\mathbb{T}$ with variability bounded by k, δ is denoted by $\mathcal{B}\Sigma\mathbb{T}_{k,\delta}$. With the notation introduced above, $\mathcal{B}\Sigma\mathbb{T}_{k,\delta} = \{b \in \mathcal{B}\Sigma\mathbb{T} \mid \forall t \in \mathbb{T} : |[t, t+\delta] \cap \tau(b)| \leq k\}$.

Similarly, a *word* $w \in \mathcal{T}\Sigma\mathbb{T}^\omega \cup \mathcal{T}\Sigma\mathbb{T}^*$ has *variability bounded* by k, δ iff for every closed interval of size δ there are at most k elements in w whose timestamps are within the interval. The set of all infinite (resp. finite) words with variability bounded by k, δ is denoted by $\mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta}$ (resp. $\mathcal{T}\Sigma\mathbb{T}^*_{k,\delta}$). With the notation introduced above, $\mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta} = \{w \in \mathcal{T}\Sigma\mathbb{T}^\omega \mid \forall i \in \mathbb{N} : t_{i+k} - t_i \geq \delta\}$ and $\mathcal{T}\Sigma\mathbb{T}^*_{k,\delta} = \{w \in \mathcal{T}\Sigma\mathbb{T}^\omega \mid \forall 0 \leq i \leq |w| - (k+1) : t_{i+k} - t_i \geq \delta\}$.

We also introduce the set of all behaviors (resp. infinite words, finite words) that are of *bounded variability for some* k, δ as $\mathcal{B}\Sigma\mathbb{T}_{\exists k \exists \delta} = \bigcup_{\substack{k \in \mathbb{N}_{>0} \\ \delta \in \mathbb{R}_{>0}}} \mathcal{B}\Sigma\mathbb{T}_{k,\delta}$ (resp. $\mathcal{T}\Sigma\mathbb{T}^\omega_{\exists k \exists \delta} = \bigcup_{\substack{k \in \mathbb{N}_{>0} \\ \delta \in \mathbb{R}_{>0}}} \mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta}$, $\mathcal{T}\Sigma\mathbb{T}^*_{\exists k \exists \delta} = \bigcup_{\substack{k \in \mathbb{N}_{>0} \\ \delta \in \mathbb{R}_{>0}}} \mathcal{T}\Sigma\mathbb{T}^*_{k,\delta}$).

2.3.2 Non-Berkeleyness

A *behavior* $b \in \mathcal{B}\Sigma\mathbb{T}$ is *non-Berkeley* for $\delta \in \mathbb{R}_{>0}$ iff every maximal constancy interval contains a closed interval of size δ. The set of all behaviors in $\mathcal{B}\Sigma\mathbb{T}$ that are non-Berkeley for δ is denoted by $\mathcal{B}\Sigma\mathbb{T}_\delta$; with the notation introduced above $\mathcal{B}\Sigma\mathbb{T}_\delta = \{b \in \mathcal{B}\Sigma\mathbb{T} \mid \forall I \in \iota(b) : \exists t \in I : [t, t+\delta] \subseteq I\}$.

Similarly, the set of infinite (resp. finite) *words* that are *non-Berkeley* for $\delta \in \mathbb{R}_{>0}$ is denoted by $\mathcal{T}\Sigma\mathbb{T}^\omega_\delta$ (resp. $\mathcal{T}\Sigma\mathbb{T}^*_\delta$) and is defined as $\mathcal{T}\Sigma\mathbb{T}^\omega_\delta = \{w \in \mathcal{T}\Sigma\mathbb{T}^\omega \mid \forall i \in \mathbb{N} : t_{i+1} - t_i \geq \delta\}$ (resp. $\mathcal{T}\Sigma\mathbb{T}^*_\delta = \{w \in \mathcal{T}\Sigma\mathbb{T}^* \mid \forall 0 \leq i \leq |w| - 2 : t_{i+1} - t_i \geq \delta\}$).

We also introduce the set of all behaviors (resp. infinite words, finite words) that are *non-Berkeley for some* $\delta \in \mathbb{R}_{>0}$ as $\mathcal{B}\Sigma\mathbb{T}_{\exists\delta} = \bigcup_{\delta \in \mathbb{R}_{>0}} \mathcal{B}\Sigma\mathbb{T}_\delta$ (resp. $\mathcal{T}\Sigma\mathbb{T}^\omega_{\exists\delta} = \bigcup_{\delta \in \mathbb{R}_{>0}} \mathcal{T}\Sigma\mathbb{T}^\omega_\delta$, $\mathcal{T}\Sigma\mathbb{T}^*_{\exists\delta} = \bigcup_{\delta \in \mathbb{R}_{>0}} \mathcal{T}\Sigma\mathbb{T}^*_\delta$).

2.3.3 Relations Among Classes

It is apparent that some of the various classes of behaviors that we introduced above are closely related. More precisely, the following inclusion relations hold.

Proposition 1. *For all $\delta' > \delta > 0$ and $k > k' \geq 2$:*

$$\mathcal{BET}_{1,\delta'} \subset \mathcal{BET}_{\delta'} \subset \mathcal{BET}_{\delta} \subset \mathcal{BET}_{k',\delta} \subset \mathcal{BET}_{k,\delta} \subset \mathcal{BET}_{\exists k \exists \delta} \subset \mathcal{BET} \tag{1}$$

$$\mathcal{BET}_{\delta} \subset \mathcal{BET}_{\exists \delta} \subset \mathcal{BET}_{\exists k \exists \delta} \subset \mathcal{BET} \tag{2}$$

$$\mathcal{BET}_{\exists \delta} \textit{ and } \mathcal{BET}_{k',\delta} \textit{ are incomparable} \tag{3}$$

$$\mathcal{TET}^{\omega}_{\delta'} \subset \mathcal{TET}^{\omega}_{\delta} = \mathcal{TET}^{\omega}_{1,\delta} \subset \mathcal{TET}^{\omega}_{k',\delta} \subset \mathcal{TET}^{\omega}_{k,\delta} \subset \mathcal{BET}_{\exists k \exists \delta} \subset \mathcal{TET}^{\omega} \tag{4}$$

$$\mathcal{TET}^{\omega}_{\delta} \subset \mathcal{TET}^{\omega}_{\exists \delta} \subset \mathcal{BET}_{\exists k \exists \delta} \subset \mathcal{TET}^{\omega} \tag{5}$$

$$\mathcal{TET}^{\omega}_{\exists \delta} \textit{ and } \mathcal{TET}^{\omega}_{k',\delta} \textit{ are incomparable} \tag{6}$$

$$\mathcal{TET}^{*}_{\delta'} \subset \mathcal{TET}^{*}_{\delta} = \mathcal{TET}^{*}_{1,\delta} \subset \mathcal{TET}^{*}_{k',\delta} \subset \mathcal{TET}^{*}_{k,\delta} \subset \mathcal{TET}^{*}_{\exists \delta} = \mathcal{TET}^{*}_{\exists k \exists \delta} = \mathcal{TET}^{*} \tag{7}$$

Proof. We prove inclusions in (1) from left to right.
$\mathcal{BET}_{1,\delta'} \subseteq \mathcal{BET}_{\delta'}$: consider any $b \in \mathcal{BET}_{1,\delta'}$ and let $t_1, t_2 \in \tau(b)$ be any two consecutive transition points in b. Then $t_2 > t_1 + \delta'$ and let $d = \min(t_2 - (t_1 + \delta'), \delta')$, thus $[t_1 + d/2, t_1 + d/2 + \delta'] \subset [t_1, t_2]$. $\mathcal{BET}_{1,\delta'} \not\supseteq \mathcal{BET}_{\delta'}$: consider b such that $\iota(b) \ni I_h = [t, t + \delta']$ for some h, t, and $|I_j| > \delta'$ for all $j \neq h$. Clearly $b \in \mathcal{BET}_{\delta'}$ but $|I_h \cap \tau(b)| = 2$, thus $b \notin \mathcal{BET}_{1,\delta'}$.
$\mathcal{BET}_{\delta'} \subseteq \mathcal{BET}_{\delta}$: immediate from the definitions as $[t, t + \delta] \subset [t, t + \delta'] \subseteq I$. $\mathcal{BET}_{\delta'} \not\supseteq \mathcal{BET}_{\delta}$: immediate for any behavior b with some $I \in \iota(b)$ such that $\delta < |I| < \delta'$.
$\mathcal{BET}_{\delta} \subseteq \mathcal{BET}_{k',\delta}$: for any $b \in \mathcal{BET}_{\delta}$ we have that for any pair of consecutive transition points $t_1, t_2 \in \tau(b)$ it is $t_2 \geq t_1 + \delta$. Then any interval $[t, t+\delta]$ contains at most 2 transition points, thus $b \in \mathcal{BET}_{k',\delta}$. $\mathcal{BET}_{\delta} \not\supseteq \mathcal{BET}_{k',\delta}$: consider b such that $t_1, t_2 \in \tau(b)$ are equal, and, for all other transition points $t_j, t_k \in \tau(b)$ it is $|t_k - t_j| > \delta$, $|t_k - t_1| > \delta$, and $|t_j - t_1| > \delta$. Clearly, $b \in \mathcal{BET}_{k',\delta}$ but $b \notin \mathcal{BET}_{\delta}$ because the interval $I = [t_1, t_2] \in \iota(b)$ is singular.
$\mathcal{BET}_{k',\delta} \subseteq \mathcal{BET}_{k,\delta}$: immediate from the definitions as $k > k'$. $\mathcal{BET}_{k',\delta} \not\supseteq \mathcal{BET}_{k,\delta}$: immediate for any behavior b with some $t_1 \leq t_2 \leq \cdots \leq t_{k+1} \in \tau(b)$ such that $t_{k+1} - t_1 = \delta$.
$\mathcal{BET}_{k,\delta} \subseteq \mathcal{BET}_{\exists k \exists \delta}$: obvious from the definitions. $\mathcal{BET}_{k,\delta} \not\supseteq \mathcal{BET}_{\exists k \exists \delta}$: immediate for any behavior b such that there exists $t_1 \leq t_2 \leq \cdots \leq t_{k+1} \in \tau(b)$ such that $t_{k+1} - t_1 < \delta$.
$\mathcal{BET}_{\exists k \exists \delta} \subseteq \mathcal{BET}$: obvious from the definitions. $\mathcal{BET}_{\exists k \exists \delta} \not\supseteq \mathcal{BET}$: consider b such that there exists a denumerable sequence of intervals $J_1, J_2, J_3, \ldots$ such that $|J_j| = |J_{j+1}|$ and $1 + |J_j \cap \tau(b)| = |J_{j+1} \cap \tau(b)|$ for all $j \geq 1$. Then $b \notin \mathcal{BET}_{k,\delta}$ for any k, δ, thus $b \notin \mathcal{BET}_{\exists k \exists \delta}$.

We prove inclusions in (2) from left to right.
$\mathcal{BET}_{\delta} \subset \mathcal{BET}_{\exists \delta}$: immediate from $\mathcal{BET}_{\delta} \subset \mathcal{BET}_{\delta_1}$ for $\delta_1 < \delta$ from (1).
$\mathcal{BET}_{\exists \delta} \subseteq \mathcal{BET}_{\exists k \exists \delta}$: immediate from $\mathcal{BET}_{\delta} \subset \mathcal{BET}_{k',\delta}$ in (1). $\mathcal{BET}_{\exists \delta} \not\supseteq \mathcal{BET}_{\exists k \exists \delta}$: immediate from $\mathcal{BET}_{\exists \delta} \not\supseteq \mathcal{BET}_{k',\delta}$ in (3).
$\mathcal{BET}_{\exists k \exists \delta} \subset \mathcal{BET}$: see (1).

We prove (3). First of all, notice that $\mathcal{BET}_{\exists \delta} \cap \mathcal{BET}_{k',\delta} \neq \emptyset$, because any $b \in \mathcal{BET}_{\delta}$ is both in $\mathcal{BET}_{k',\delta}$ from (1) and in $\mathcal{BET}_{\exists \delta}$ from (2).
$\mathcal{BET}_{\exists \delta} \not\subseteq \mathcal{BET}_{k',\delta}$: consider b such that any $I \in \iota(b)$ is such that $|I| \geq \delta/(2k')$ and there exist $k' + 1$ transition points $t_1 \leq t_2 \leq \ldots \leq t_{k'+1} \in \iota(b)$ such that $t_{k'+1} - t_1 \leq \delta$. Hence $b \in \mathcal{BET}_{\delta'}$ for $\delta' < \delta/(2k')$ and thus $b \in \mathcal{BET}_{\exists \delta}$, but $b \notin \mathcal{BET}_{k',\delta}$.
$\mathcal{BET}_{\exists \delta} \not\supseteq \mathcal{BET}_{k',\delta}$: consider b such that there exists exactly one singular interval $I = [t, t] \in \iota(b)$ and for all other intervals I' it is $|I'| > \delta$. Then, $b \in \mathcal{BET}_{2,\delta}$ and

thus also $\mathcal{B}\Sigma\mathbb{T}_{k',\delta}$ from (1). However, $b \notin \mathcal{B}\Sigma\mathbb{T}_{\exists\delta}$ because of the singular interval I.

We prove inclusions in (4) from left to right.
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\delta'} \subset \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\delta}$: obvious from the definitions.
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\delta} = \mathcal{T}\Sigma\mathbb{T}^{\omega}_{1,\delta}$: also obvious from the definitions.
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{1,\delta} \subseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta}$: immediate for $t_{i+1} - t_i \geq \delta$ implies $t_{i+k'} - t_i \geq \delta$ for all $k' \geq 1$. $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{1,\delta} \not\supseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta}$: immediate for any word w with some $t_1 \leq t_2 \leq t_3$ such that $t_3 - t_1 = \delta$.
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta} \subseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k,\delta}$: immediate for $t_{i+k'} - t_i \geq \delta$ implies $t_{i+k} - t_i \geq \delta$ for all $k > k'$.
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta} \not\supseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k,\delta}$: immediate for any word w with some $t_1 \leq t_2 \leq \cdots \leq t_{k+1}$ such that $t_{k+1} - t_1 = \delta$.
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{k,\delta} \subseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists k \exists \delta}$: obvious from the definitions. $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{k,\delta} \not\supseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists k \exists \delta}$: immediate for any word w with some $t_1 \leq t_2 \leq \cdots \leq t_{k+1} \leq t_{k+2}$ such that $t_{k+2} = \delta$; clearly $w \notin \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k,\delta}$.
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists k \exists \delta} \subseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}$: obvious from the definitions. $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists k \exists \delta} \not\supseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}$: let w be a word such that there exist $t_{2^j} \leq t_{2^j+1} \leq \cdots \leq t_{2^j+j}$ such that $t_{2^j+j} - t_{2^j} = t_{2^{j+1}+j+1} - t_{2^{j+1}}$ for all $j \geq 1$. Clearly $w \notin \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k,\delta}$ for any k, δ.

We prove inclusions in (5) from left to right.
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\delta} \subseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta}$: obvious from the definitions. $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\delta} \not\supseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta}$: immediate from $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\delta} \subset \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\delta_1}$ for $\delta_1 < \delta$ from (4).
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta} \subseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists k \exists \delta}$: immediate from $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\delta} \subset \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta}$ in (4). $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta} \not\supseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists k \exists \delta}$: immediate from $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta} \not\supseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta}$ in (6).
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists k \exists \delta} \subset \mathcal{T}\Sigma\mathbb{T}^{\omega}$: see (4).

We prove (6). First of all, notice that $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta} \cap \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta} \neq \emptyset$, because any $w \in \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\delta}$ is both in $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta}$ from (4) and in $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta}$ from (5).
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta} \not\subseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta}$: consider w such that $t_{i+1} - t_i = \delta/(k'+1)$ for all $i \in \mathbb{N}$. Then, $w \in \mathcal{T}\Sigma\mathbb{T}^{\omega}\delta'$ for $\delta' = \delta/(k'+1)$, so $w \in \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta}$. However, $t_{i+k'} - t_i = k'\delta/(k'+1) < \delta$ thus $w \notin \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta}$.
$\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta} \not\supseteq \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta}$: consider word w such that $t_{2i} = 2i\delta$ and $t_{2i+1} = t_{2(i+1)} - 1/2^i$ for all $i \in \mathbb{N}$. Then, $w \notin \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\delta}$ for any $\delta > 0$, hence $w \notin \mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists\delta}$. However, $w \in \mathcal{T}\Sigma\mathbb{T}^{\omega}_{2,\delta}$, hence $w \in \mathcal{T}\Sigma\mathbb{T}^{\omega}_{k',\delta}$ for all $k' \geq 2$ from (4).

We prove inclusions in (7) from right to left. We only consider inclusions that are different from the analogous ones for infinite timed words, as the other are the same as in (4) or (5).
$\mathcal{T}\Sigma\mathbb{T}^{*}_{\exists\delta} \subseteq \mathcal{T}\Sigma\mathbb{T}^{*}_{\exists k \exists \delta} \subseteq \mathcal{T}\Sigma\mathbb{T}^{*}$: obvious from the definitions. $\mathcal{T}\Sigma\mathbb{T}^{*}_{\exists\delta} \supseteq \mathcal{T}\Sigma\mathbb{T}^{*}_{\exists k \exists \delta} \supseteq \mathcal{T}\Sigma\mathbb{T}^{*}$: let $w \in \mathcal{T}\Sigma\mathbb{T}^{*}$; since $|w|$ is finite, $m = \min_{i \in \mathbb{N}}(t_{i+1} - t_i)$ is also finite (and greater than zero because we assume strictly monotonic timestamps). Hence, $w \in \mathcal{T}\Sigma\mathbb{T}^{*}_{m} = \mathcal{T}\Sigma\mathbb{T}^{*}_{1,m}$ and thus $w \in \mathcal{T}\Sigma\mathbb{T}^{*}_{\exists\delta}$ and $w \in \mathcal{T}\Sigma\mathbb{T}^{*}_{\exists k \exists \delta}$.
$\mathcal{T}\Sigma\mathbb{T}^{*}_{k,\delta} \subseteq \mathcal{T}\Sigma\mathbb{T}^{*}_{\exists\delta}$: obvious because $\mathcal{T}\Sigma\mathbb{T}^{*}_{k,\delta} \subseteq \mathcal{T}\Sigma\mathbb{T}^{*}$ by definition and we proved that $\mathcal{T}\Sigma\mathbb{T}^{*} = \mathcal{T}\Sigma\mathbb{T}^{*}_{\exists\delta}$. $\mathcal{T}\Sigma\mathbb{T}^{*}_{k,\delta} \not\supseteq \mathcal{T}\Sigma\mathbb{T}^{*}_{\exists\delta}$: consider w such that there exist $k+1$ consecutive timestamps $t_0 < t_1 < \cdots < t_{k+1}$ such that $t_{k+1} - t_0 = \delta$ and for all timestamps t_j with $j > k+1$, $t_j - t_{k+1} > \delta$. Clearly, $w \in \mathcal{T}\Sigma\mathbb{T}^{*}_{m}$ for $m = \min_{i \in [0..k]}(t_{i+1} - t_i)$ so $w \in \mathcal{T}\Sigma\mathbb{T}^{*}_{\exists\delta}$, but $w \notin \mathcal{T}\Sigma\mathbb{T}^{*}_{k,\delta}$. □

Remark 2. In the rest of the paper we will consider only behaviors (and words) that have bounded variability δ for some *rational* value of $\delta > 0$. This is due to the fact that even decidable logics such as MITL become undecidable if irrational constants are allowed [Rab07]. It is also well-known that this is without loss of generality — as much as satisfiability is concerned — because formulas of common temporal logics are satisfiable iff they are satisfiable on behaviors (or words) with rational transition points [AFH96]. Finally, it is clear from Proposition 1 that, for any irrational Δ, we can pick a rational δ that approximates Δ with an arbitrary precision and such that the semantic class for Δ is strictly contained in the corresponding class for δ.

3 MTL and Its Relatives

The main focus of this paper is the decidability of MTL over the classes of behaviors and words that we introduced in the previous section. Hence, this section introduces formally MTL and other closely related temporal logics that will be used to obtain the results of the following sections. For notational convenience, in this paper we usually denote MITL formulas as ψ and MTL formulas as ϕ.

3.1 MITL and MTL

MITL. Let us start with the Metric Interval Temporal Logic (MITL) [AFH96], a decidable subset of MTL. MITL formulas are defined as follows, for $\mathsf{p} \in \mathcal{P}$ an atomic proposition and I a non-singular interval of the nonnegative reals with rational (or unbounded) endpoints.

$$\text{MITL} \ni \psi := \mathsf{p} \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \mathsf{U}_I(\psi_1, \psi_2) \mid \mathsf{S}_I(\psi_1, \psi_2)$$

MTL. Metric Temporal Logic (MTL) [AH93] is defined simply as an extension of MITL where singular intervals are allowed.

Derived operators. Abbreviations such as $\top, \bot, \vee, \Rightarrow, \Leftrightarrow$ are defined as usual. We drop the interval I in operators when it is $(0, \infty)$, and we represent intervals by pseudo-arithmetic expressions such as $> k$, $\geq k$, $< k, \leq k$, and $= k$ for (k, ∞), $[k, \infty)$, $(0, k)$, $(0, k]$ and $[k, k]$, respectively.

We also introduce a few derived temporal operators; in the following definitions I is an interval of the nonnegative reals with rational (or unbounded) endpoints. More precisely, the following definitions introduce MITL derived operators if I is taken to be non-singular and ϕ is an MITL formula; otherwise they introduce MTL derived operators. For both semantics we introduce the following derived operators: $\Diamond_I(\phi) = \mathsf{U}_I(\top, \phi)$, $\Box_I(\phi) = \neg\Diamond_I(\neg\phi)$, $\mathsf{R}_I(\phi_1, \phi_2) = \neg\mathsf{U}_I(\neg\phi_1, \neg\phi_2)$, as well as their past counterparts $\overleftarrow{\Diamond}_I(\phi) = \mathsf{S}_I(\top, \phi)$, $\overleftarrow{\Box}_I(\phi) = \neg\overleftarrow{\Diamond}_I(\neg\phi)$, $\mathsf{T}_I(\phi_1, \phi_2) = \neg\mathsf{S}_I(\neg\phi_1, \neg\phi_2)$, and $\text{Alw}(\phi) = \overleftarrow{\Box}(\phi) \wedge \phi \wedge \Box(\phi)$. Just for the behavior semantics we also introduce $\bigcirc(\phi) = \mathsf{U}(\phi, \top)$, $\overleftarrow{\bigcirc}(\phi) = \mathsf{S}(\phi, \top)$, and $\triangle(\phi) = \overleftarrow{\bigcirc}(\neg\phi) \wedge (\phi \vee \bigcirc(\phi))$. Just for the word semantics we have instead $\bigcirc_I(\phi) = \mathsf{U}_I(\bot, \phi)$ and $\overleftarrow{\bigcirc}_I(\phi) = \mathsf{S}_I(\bot, \phi)$.

Granularity. The granularity ρ of an M[I]TL formula ϕ is defined as the reciprocal of the product of all positive finite denominators appearing in intervals of formulas. Informally, the granularity characterizes the "resolution" of the formula: any change in a behavior by an amount smaller than ρ cannot be detected by a formula whose granularity is ρ. Let MITL^ρ and $\text{MITL}^{\geq\rho}$ (resp. MTL^ρ and $\text{MTL}^{\geq\rho}$) denote the sets of all MITL (resp. MTL) formulas of granularity ρ and greater than or equal to ρ, respectively.

Semantics. For $b \in \overline{\mathcal{B}\Sigma\mathbb{T}}$ (with $\Sigma = 2^{\mathcal{P}}$) and $t \in \mathbb{T}$ the semantics of MTL (and MITL) is defined as follows.[2]

$b(t) \models_{\mathbb{T}} \mathsf{p}$	iff	$\mathsf{p} \in b(t)$
$b(t) \models_{\mathbb{T}} \neg\phi$	iff	$b(t) \not\models_{\mathbb{T}} \phi$
$b(t) \models_{\mathbb{T}} \phi_1 \wedge \phi_2$	iff	$b(t) \models_{\mathbb{T}} \phi_1$ and $b(t) \models_{\mathbb{T}} \phi_2$
$b(t) \models_{\mathbb{T}} \mathsf{U}_I(\phi_1, \phi_2)$	iff	there exists $d \in t \oplus I \cap \mathbb{T}$ such that $b(d) \models_{\mathbb{T}} \phi_2$ and for all $u \in (t, d)$ it is $b(u) \models_{\mathbb{T}} \phi_1$
$b(t) \models_{\mathbb{T}} \mathsf{S}_I(\phi_1, \phi_2)$	iff	there exists $d \in -I \oplus t \cap \mathbb{T}$ such that $b(d) \models_{\mathbb{T}} \phi_2$ and for all $u \in (d, t)$ it is $b(u) \models_{\mathbb{T}} \phi_1$
$b \models_{\mathbb{T}} \phi$	iff	$b(0) \models_{\mathbb{T}} \phi$

For $w = (\sigma, \mathbf{t}) \in \overline{\mathcal{T}\Sigma\mathbb{T}^\omega \cup \mathcal{T}\Sigma\mathbb{T}^*}$ (with $\Sigma = 2^{\mathcal{P}}$) and $\mathbb{N} \ni k < |w|$ the semantics of MTL (and MITL) is defined as follows.

$w(k) \models_{\mathbb{T}} \mathsf{p}$	iff	$\mathsf{p} \in \sigma_k$
$w(k) \models_{\mathbb{T}} \neg\phi$	iff	$w(k) \not\models_{\mathbb{T}} \phi$
$w(k) \models_{\mathbb{T}} \phi_1 \wedge \phi_2$	iff	$w(k) \models_{\mathbb{T}} \phi_1$ and $w(k) \models_{\mathbb{T}} \phi_2$
$w(k) \models_{\mathbb{T}} \mathsf{U}_I(\phi_1, \phi_2)$	iff	there exists $k < h < |w|$ such that $t_h \in I \oplus t_k$, $w(h) \models_{\mathbb{T}} \phi_2$, and for all $k < j < h$ it is $w(j) \models_{\mathbb{T}} \phi_1$
$w(k) \models_{\mathbb{T}} \mathsf{S}_I(\phi_1, \phi_2)$	iff	there exists $0 \leq h < k$ such that $t_h \in -I \oplus t_k$, $w(h) \models_{\mathbb{T}} \phi_2$, and for all $h < j < k$ it is $w(j) \models_{\mathbb{T}} \phi_1$
$w \models_{\mathbb{T}} \phi$	iff	$w(0) \models_{\mathbb{T}} \phi$

Normal form over behaviors. In order to simplify the presentation of some of the following results, we present a normal form for MITL over behaviors, defined by the following grammar, where d is a positive rational number.

$$\psi \ := \ \mathsf{p} \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \mathsf{U}(\psi_1, \psi_2) \mid \mathsf{S}(\psi_1, \psi_2) \mid \Diamond_{<\mathrm{d}}(\psi) \mid \overleftarrow{\Diamond}_{<\mathrm{d}}(\psi)$$

The fact that every MITL formula can be expressed according to the syntax above follows from two results. [HR04, Th. 4.1, Prop. 4.2] showed that every generic MITL formula using intervals with integer endpoints can be translated into an equivalent one in the normal form above. Second, [AFH96, Lm. 2.16] showed that every MITL using intervals with rational endpoints can be translated into an equi-satisfiable one with integer endpoints only; this is achieved by uniformly scaling the endpoints into integers. It is then clear that all our results for behaviors can assume formulas in this normal form.

In addition, an analogous normal form for MTL can be defined by introducing the additional operators:

$$\phi \ := \ \Diamond_{=\mathrm{d}}(\phi) \mid \overleftarrow{\Diamond}_{=\mathrm{d}}(\phi)$$

Derived behavior of a formula. For any MTL formula ϕ and behavior $b \in \mathcal{B}\Sigma\mathbb{T}$, we define the derived behavior b_ϕ that represents the truth value of ϕ over b; namely:

$$b_\phi(t) = \begin{cases} b(t) \cup \{\phi\} & \text{if } b(t) \models_{\mathbb{T}} \phi \\ b(t) & \text{otherwise} \end{cases}$$

[2]We assume that $0 \in \mathbb{T}$ without practical loss of generality.

Size of a formula. The size $|\phi|$ of a formula ϕ is defined as the number of its atomic propositions, connectives, and temporal operators, multiplied by the size — assuming a binary encoding — of the largest finite constant appearing in intervals bounding temporal operators.

3.2 QTL and QITL: Decidable Extensions of MITL

Following [HR04, Rab08], we introduce extensions of MITL over behaviors that are known to be decidable. They will be useful in the decidability proofs of Section 5.

QTL(n). For $n \in \mathbb{N}_{>0}, \mathrm{d} \in \mathbb{Q}_{\geq 0}, \psi_1, \ldots, \psi_n \in$ MITL, we introduce the n-ary modality $\diamondsuit^{n}_{<\mathrm{d}}(\psi_1, \ldots, \psi_n)$.[3] For every $n \in \mathbb{N}_{>0}$, we denote by QTL(n) the temporal logic obtained by extending MITL with all operators $\diamondsuit^{1}_{<\mathrm{d}}(), \ldots, \diamondsuit^{n}_{<\mathrm{d}}()$.

Informally, $\diamondsuit^{n}_{<\mathrm{d}}(\psi_1, \ldots, \psi_n)$ specifies that the formulas $\psi_1, \psi_2, \ldots, \psi_n$ will occur within d time units, in that order. Formally, we have the following semantics over behaviors:

$b(t) \models_{\mathbb{T}} \diamondsuit^{n}_{<\mathrm{d}}(\psi_1, \ldots, \psi_n)$ iff there exist $t < t_1 < \cdots < t_n < t + \mathrm{d}$ such that for all $1 \leq i \leq n$ it is $b(t_i) \models_{\mathbb{T}} \psi_i$

Obviously, QTL(1) is exactly MITL. Then, we recall the following from [HR04, Th. 10.2].

Proposition 3 (Expressiveness and Decidability of QTL(n))**.** *For all $n > 0$: (1)* QTL(n) *is decidable; and (2)* QTL($n+1$) *is strictly more expressive than* QTL(n)*.*

QITL(n). We further extend QTL by introducing modalities $\diamondsuit^{n}_{I}(\psi_1, \ldots, \psi_n)$ for $n > 0$ and I a non-singular interval. We denote the corresponding temporal logics by QITL(n), for $n > 0$. Also, we denote the temporal logic $\bigcup_{k>0}$ QITL(k) simply by QITL. The semantics of the new operators is as expected:

$b(t) \models_{\mathbb{T}} \diamondsuit^{n}_{I}(\psi_1, \ldots, \psi_n)$ iff there exist $t_1 < \cdots < t_n \in I \oplus t$ such that for all $1 \leq i \leq n$ it is $b(t_i) \models_{\mathbb{T}} \psi_i$

We also introduce the abbreviation $\diamondsuit^{n}_{I}(\psi) = \diamondsuit^{n}_{I}\left(\underbrace{\psi, \psi, \ldots, \psi}_{n \text{ times}}\right)$.

Note that QITL is essentially equivalent to the logic TPLI introduced in [Rab08]. More precisely, the fundamental difference between the two logics is that TLPI allows only open intervals (l, u) in the $\diamondsuit^{n}_{(l,u)}(\cdots)$ operators. It is however clear that this gap can be bridged along the lines of [HR04]. Namely, the

[3] In [HR04] the same modality is denoted as $(\Diamond\psi_1 \cdots \Diamond\psi_n)_{\mathrm{d}}$.

following equivalences hold, where the modality $\mathrm{Pn}_n^{(l,u)}$ is defined in [Rab08].[4]

$$\diamondsuit^{n}_{(l,u)}(\psi_1,\ldots,\psi_n) \equiv \mathrm{Pn}_n^{(l,u)}(\psi_1,\ldots,\psi_n)$$

$$\diamondsuit^{n}_{[l,u)}(\psi_1,\ldots,\psi_n) \equiv \diamondsuit^{n}_{(l,u)}(\psi_1,\ldots,\psi_n) \vee \left(\begin{array}{c} \diamondsuit^{n-1}_{(l,u)}(\psi_2 \wedge \mathsf{S}(\neg\psi_1,\psi_1), \psi_3,\ldots,\psi_n) \\ \wedge \overleftarrow{\bigcirc}\Big(\diamondsuit^{n}_{(l,u)}(\psi_1,\ldots,\psi_n)\Big) \end{array} \right)$$

$$\diamondsuit^{n}_{(l,u]}(\psi_1,\ldots,\psi_n) \equiv \diamondsuit^{n}_{(l,u)}(\psi_1,\ldots,\psi_n) \vee \left(\begin{array}{c} \diamondsuit^{n-1}_{(l,u)}(\psi_1,\ldots,\psi_{n-1} \wedge \mathsf{U}(\neg\psi_n,\psi_n)) \\ \wedge \bigcirc\Big(\diamondsuit^{n}_{(l,u)}(\psi_1,\ldots,\psi_n)\Big) \end{array} \right)$$

$$\diamondsuit^{n}_{[l,u]}(\psi_1,\ldots,\psi_n) \equiv \diamondsuit^{n-1}_{[l,u)}(\psi_1,\ldots,\psi_{n-1}) \wedge \diamondsuit^{n-1}_{(l,u]}(\psi_2,\ldots,\psi_n)$$

Hence, the following is a corollary of the complexity results for TLPI presented in [Rab08].

Proposition 4 (Decidability and Complexity of QITL)**.** QITL *is decidable with a* **PSPACE***-complete validity problem (***EXPSPACE***-complete if constants are encoded succinctly).*

As a final remark, let us show that the results of this paper are even slightly more robust, as they can be proved with the following *weakening* of Proposition 4. Let $\mathrm{QITL}^{<B}$ denote the set of QITL formulas where all intervals I appearing in $\diamondsuit^{n}_{I}(\cdots)$ operators (for any n) have size bounded by B.

Proposition 5. *For any* $B \in \mathbb{N}_{>0}$*,* $\mathrm{QITL}^{<B}$ *is decidable with an* **EXPSPACE***-complete if constants are encoded succinctly.*

Proposition 5 can be proved through the following.

Lemma 6. *Any* QITL *formula* ψ *can be translated into a TPL [Rab08] formula* ψ' *such that* $|\psi'| = |\psi|^{\mathrm{O}(L)}$*, where* L *is the size of the largest (finite) interval used in* ψ*.*

Proof. Without practical loss of generality, assume $l, u \in \mathbb{N}$. Consider the following equivalences, for $1 < l - u \leq L$:

$$\diamondsuit^{0}_{I}(\cdots) \equiv \top$$

$$\diamondsuit^{n}_{(0,1)}(\psi_1,\ldots,\psi_n) \equiv \mathrm{Pn}_n(\psi_1,\ldots,\psi_n)$$

$$\diamondsuit^{n}_{(l,l+1)}(\psi_1,\ldots,\psi_n) \equiv \Box_{(0,1)}\Big(\Diamond_{(0,l)}\Big(\diamondsuit^{n}_{(0,1)}(\psi_1,\ldots,\psi_n)\Big)\Big)$$

$$\diamondsuit^{n}_{(l,u)}(\psi_1,\ldots,\psi_n) \equiv \bigvee_{\substack{n=\Sigma_{i=1}^{u-l} k_i \\ k_i \in \mathbb{N}}} \left(\begin{array}{c} \bigwedge_{i=1}^{u-l-1} \diamondsuit^{k_i}_{(l+i-1,l+i]} \left(\begin{array}{c} \psi_{1+\Sigma_{j<i} k_j} \\ ,\ldots, \\ \psi_{k_i+\Sigma_{j<i} k_j} \end{array} \right) \\ \\ \wedge\ \diamondsuit^{k_{u-l}}_{(u-1,u)}\left(\psi_{n-k_{u-l}+1},\ldots,\psi_n\right) \end{array} \right)$$

where closed (or half-closed) intervals can be handled as discussed above. Essentially, to represent intervals of size larger than one, we consider all possible ways in which the n formulas can be distributed among the $u - l$ adjacent intervals of unit length, and we explicitly enumerate them. From the definitions for TPL in [Rab08], it should be clear that the truth-preserving translation induced by

[4]For simplicity, we assume l, u to be natural numbers.

the equivalences builds TPL formulas whose size is as described in the statement of the lemma; in particular, the number of terms in the last disjunction is upper-bounded by $n^{u-l} \leq n^L$. □

Correspondingly, Proposition 5 follows from Lemma 6 and the fact — also proved in [Rab08] — that TPL is **EXPSPACE**-complete. In this paper, we are going to use the $\diamondsuit_I^n(\cdots)$ operator only with intervals of size at most δ, where δ will be independent of the size of the other formulas. Hence, Proposition 5 suffices to prove all the decidability results of this paper.

4 Syntactic Definition of Regularity Constraints

In this section we show how to express the regularity constraints of bounded variability and non-Berkeleyness as MITL or QITL formulas. The following two sub-sections introduce two preliminary results.

4.1 From Non-Berkeleyness to Bounded Variability

Let ϕ be any MTL formula and $b \in \mathcal{B}\Sigma\mathbb{T}_\delta$ a non-Berkeley behavior. While non-Berkeleyness is defined according to the behavior of atomic propositions in b, it is simple to realize that, in general, it cannot be lifted to the behavior of ϕ itself in b_ϕ. In other words, it may happen that b_ϕ is Berkeley (i.e., two adjacent transitions are less than δ time units apart) even if b is not.

However, b_ϕ is at least with variability bounded by $\theta(\phi), \delta$, where $\theta(\phi)$ can be computed from the structure of ϕ. More precisely, consider the following definition, where β is a Boolean combination of atomic propositions.

$$\begin{aligned}
\theta(\beta) &= 2 \\
\theta(\neg\phi) &= \theta(\phi) \\
\theta(\phi_1 \wedge \phi_2) &= \theta(\phi_1) + \theta(\phi_2) \\
\theta(\mathsf{U}(\phi_1, \phi_2)) &= \theta(\phi_1) \\
\theta(\diamondsuit_{<\mathrm{d}}(\phi)) &= \theta(\phi) + 1 \\
\theta(\diamondsuit_{=\mathrm{d}}(\phi)) &= \theta(\phi)
\end{aligned}$$

Note that $\theta(\phi) = \mathrm{O}(|\phi|)$.

Then, we can prove the following.

Lemma 7. *For any $b \in \mathcal{B}\Sigma\mathbb{T}_\delta$ and MTL formula ϕ, it is $b_\phi \in \mathcal{B}\Sigma\mathbb{T}_{\theta(\phi),\delta}$.*

Proof. Let b be a non-Berkeley behavior for $\delta > 0$, $J = [t, t+\delta]$ be any closed interval of size δ, and ϕ be a generic MTL formula. The proof goes by induction on the structure of ϕ.

- $\phi = \beta$.
  Since b is non-Berkeley, for any two adjacent transition points $x_i, x_{i+1} \in \tau(b_\beta) = \tau(b)$ it is $x_{i+1} - x_i \geq \delta$. Therefore, J contains at most two transitions of b_β, and in fact $\theta(\beta) = 2$.

- $\phi = \neg\phi'$.
  Clearly $\tau(b'_\phi) = \tau(b_\phi)$; therefore b_ϕ has at most $\theta(\phi) = \theta(\phi')$ transition points within J, by inductive hypothesis.

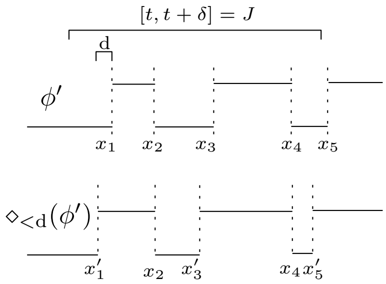

Figure 1: $b_{\phi'}$ and $b_{\Diamond_{<\mathrm{d}}(\phi')}$ over J.

- $\phi = \phi_1 \wedge \phi_2$.
  A little reasoning should convince us that $\tau(b_\phi) \subseteq \tau(b_{\phi_1}) \cup \tau(b_{\phi_2})$; in fact whenever b_ϕ has a transition, at least one of b_{ϕ_1}, b_{ϕ_2} must have a transition, while the converse does not hold necessarily. Hence, $|\tau(b_\phi) \cap J| \leq |(\tau(b_{\phi_1}) \cup \tau(b_{\phi_2})) \cap J| \leq |\tau(b_{\phi_1}) \cap J| + |\tau(b_{\phi_2}) \cap J| \leq \theta(\phi_1) + \theta(\phi_2) = \theta(\phi)$, where the last inequality follows by inductive hypothesis.

- $\phi = \Diamond_{=\mathrm{d}}(\phi')$.
  Clearly, $\tau(b_\phi) = \ldots, x_{-1} - \mathrm{d}, x_0 - \mathrm{d}, x_1 - \mathrm{d}, \ldots, x_i - \mathrm{d}, \ldots$, where $\tau(\phi') = \ldots, x_{-1}, x_0, x_1, \ldots, x_i, \ldots$. Thus, $b_{\phi'} \in \mathbb{BET}_{\theta(\phi'),\delta}$ implies $b_\phi \in \mathbb{BET}_{\theta(\phi),\delta}$ as well, since $\theta(\phi) = \theta(\phi')$.

- $\phi = \Diamond_{<\mathrm{d}}(\phi')$.
  Let $x_1, \ldots, x_k = \tau(b_{\phi'}) \cap J$ be the transition points of $b_{\phi'}$ over J; by inductive hypothesis we know that $k \leq \theta' = \theta(\phi')$.

  Let us first consider the case: $x_i - x_{i-1} \geq \mathrm{d}$ for all $i = 2, \ldots, k+1$. If also x_1 is a transition from false to true (see Figure 1 for an example with $k = 4$, where $x'_i = x_i - \mathrm{d}$), b_ϕ has the corresponding transition points $x_1 - \mathrm{d}, x_2, x_3 - \mathrm{d}, \ldots$; if instead x_1 is a transition from true to false, b_ϕ has the corresponding transition points $x_1, x_2 - \mathrm{d}, x_3, \ldots$. In particular, note that when $x_{i+1} - x_i = \mathrm{d}$ and ϕ' is false throughout (x_i, x_{i+1}), $x_i = x_{i+1} - \mathrm{d}$ is a punctual transition point for b_ϕ, and in fact it appears twice in $\tau(b_\phi)$. Overall, b_ϕ has at most all the transition points $b_{\phi'}$ has over J, plus one corresponding to $x_{k+1} - \mathrm{d}$. Since $\theta(\phi) = \theta(\phi') + 1$, we have that $b_\phi \in \mathbb{BET}_{\theta(\phi),\delta}$.

  Whenever $x_i - x_{i-1} < \mathrm{d}$ for some $i = 2, \ldots, k+1$, the transition points of $b_{\phi'}$ may instead be fewer. In fact, if x_1 is a transition from false to true, for all odd $i = 3, \ldots, k+1$ such that $x_i - x_{i-1} < \mathrm{d}$, there are no transition points for b_ϕ between x_{i-1} and x_{i+1}. Similarly, if x_1 is a transition from true to false, for all even $i = 2, \ldots, k+1$ such that $x_i - x_{i-1} < \mathrm{d}$, there are no transition points for b_ϕ between x_{i-1} and x_{i+1}. Overall, $\theta(\phi) = \theta(\phi')+1$ is an upper bound on the number of transitions of b_ϕ over J in this case as well.

- $\phi = \mathsf{U}(\phi_1, \phi_2)$.
  Let $x_1, \ldots, x_k = \tau(b_{\phi_1}) \cap J$ be the transition points of b_{ϕ_1} over J; by inductive hypothesis we know that $k \leq \theta_1 = \theta(\phi_1)$. For the sake of

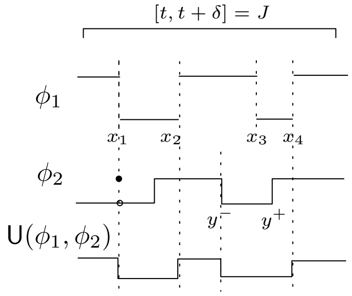

Figure 2: b_{ϕ_1}, b_{ϕ_2} and $b_{\mathsf{U}(\phi_1,\phi_2)}$ over J.

presentation, assume that x_1 is a transition from true to false; we will show shortly that the other cases can be handled with trivial modifications.

Let us first consider the case $b(x_i) \models_{\mathbb{T}} \phi_2 \vee \overleftarrow{\bigcirc}(\phi_2)$ for all odd $i = 1, \dots, k$. Then, it is clear that ϕ holds throughout $(t, x_1), (x_2, x_3), \dots, (x_{2h}, x_{2h+1})$, $\dots$, and it does not hold over $(x_1, x_2), (x_3, x_4), \dots, (x_{2h-1}, x_{2h}), \dots$. Hence, the transition points of ϕ over J are precisely $x_1, x_2, \dots, x_k$, which shows that $b_\phi \in \mathbb{BET}_{\theta(\phi),\delta}$, as $\theta(\phi) = \theta_1$.

Let us now consider the case in which some x_i, with odd i, is such that $b(x_i) \models_{\mathbb{T}} \neg\phi_2 \wedge \overleftarrow{\bigcirc}(\neg\phi_2)$. Let us denote by y^- (resp. y^+) the largest (resp. smallest) transition point of ϕ_2 which comes before (resp. after) x_i (see Figure 2 for a visual representation in which $i = 3$). Then, it should be clear that ϕ holds over (x_{i-1}, y^-) and is false over (y^-, x_{i+1}), where (x_{i-1}, y^-) can possibly be empty. So, x_i is surely not a transition point for ϕ, but y^- can be (if (x_{i-1}, y^-) is not empty). Iterating this reasoning for all odd i's such that $b(x_i) \models_{\mathbb{T}} \neg\phi_2 \wedge \overleftarrow{\bigcirc}(\neg\phi_2)$, we have that the number of transition points of b_ϕ over J is at most θ_1, so $b_\phi \in \mathbb{BET}_{\theta(\phi),\delta}$ in this case as well.

The remaining cases can be handled routinely. In particular, if x_1 is a transition from false to true we just replace "odd" with "even" in the above reasoning. □

4.2 Describing Sequences of Transitions

This section introduces QITL formulas that can be used to describe sequences of transitions of the truth value of MTL formulas.

For every QITL formula ϕ, nonsingular interval I, and integer $k > 0$, we introduce the QITL formula:

$$\mathsf{happ}(\phi, k, I) = \Diamond\!\!\!\Diamond_I^k \left(\underbrace{\phi', \neg\phi', \cdots, \phi, \neg\phi, \phi}_{k \text{ terms}} \right) \wedge \neg \Diamond\!\!\!\Diamond_I^{k+1} \left(\underbrace{\phi', \neg\phi', \cdots, \phi, \neg\phi, \phi, \neg\phi}_{k+1 \text{ terms}} \right)$$

where $\phi' = \phi$ if k is odd and $\phi' = \neg\phi$ otherwise.

Informally, $\mathsf{happ}(\phi, k, I)$ means that ϕ takes exactly $k - 1$ consecutive transitions, eventually leading to true, i.e., it holds at the end of I. More precisely we have the following.

Lemma 8. *Let ϕ undergo exactly k transitions over $t \oplus [\tau^-, \tau^+]$, for some $\tau^+ > \tau^- \geq 0$, that is $\langle \tau(b_\phi) \cap t \oplus [\tau^-, \tau^+] \rangle = k$; then, $b(t + \tau^+) \models \phi$ iff $b(t) \models \mathsf{happ}(\phi, k+1, [\tau^-, \tau^+])$.*

Proof. For the $\Rightarrow$ direction, it is clear that $\diamondsuit^{k+2}_{[\tau^-,\tau^+]}(\phi', \neg\phi', \cdots, \phi, \neg\phi, \phi, \neg\phi)$ cannot hold at t, since ϕ undergoes no more than k transitions over $t \oplus [\tau^-, \tau^+]$, and such transitions are represented by no more than $k+1$ alternations between false and true values. From the fact that ϕ holds at $t + \tau^+$, we infer that the k-th transition over the interval is to a true value, which is held until $t + \tau^+$ included. Thus ϕ alternates over a total of $k+1$ false and true values, i.e., $\diamondsuit^{k+1}_{[\tau^-,\tau^+]}(\phi', \neg\phi', \cdots, \phi, \neg\phi, \phi)$ holds at t.

For the $\Leftarrow$ direction, let us start from the fact that the k-th transition is the last one over the interval $t \oplus [\tau^-, \tau^+]$ and it yields a true value; also, the value is kept until $t + \tau^+$ included, otherwise there would be at least one more transition to false. Consequently, $b(t + \tau^+) \models \phi$. □

Next, we introduce a formula, built upon $\mathsf{happ}(\phi, k, I)$, to describe the case where we have *at most* n transitions over I. For every QITL formula ϕ, non-singular interval I, and $n > 0$, we introduce the QITL formula:

$$\mathsf{yieldsT}(\phi, n, I) = \bigvee_{0 \leq k \leq n} \mathsf{happ}(\phi, k, I) \tag{8}$$

Generalizing Lemma 8 we have the following.

Lemma 9. *Let ϕ undergo at most n transitions over $t \oplus [\tau - \delta, \tau]$ for some $\tau > 0$, that is $\langle \tau(b_\phi) \cap t \oplus [\tau - \delta, \tau] \rangle \leq n$; then $b(t + \tau) \models \phi$ iff $b(t) \models \mathsf{yieldsT}(\phi, n + 1, [\max(0, \tau - \delta), \tau])$.*

Proof. Let ϕ undergo exactly $k \leq n$ transitions over $t \oplus [\tau - \delta, \tau]$, and let $I = [\max(0, \tau - \delta), \tau]$.

Let us first consider the case $\tau - \delta > 0$, and thus $I = [\tau - \delta, \tau]$. If $b(t+\tau) \models \phi$ then, from Lemma 8, $b(t) \models \mathsf{happ}(\phi, k+1, [\tau - \delta, \tau])$, which implies $b(t) \models \mathsf{yieldsT}(\phi, n+1, [\tau - \delta, \tau])$ according to (8). Conversely, if $b(t) \models \mathsf{yieldsT}(\phi, n+1, [\tau - \delta, \tau])$ then $b(t) \models \mathsf{happ}(\phi, \widetilde{k}, [\tau - \delta, \tau])$ for some $\widetilde{k} \leq n+1$. In particular, it is $b(t) \models \mathsf{happ}(\phi, k+1, I)$; hence $b(t + \tau) \models \phi$ from Lemma 8.

Let us now assume $\tau - \delta \leq 0$, and thus $I = [0, \tau] \subseteq [\tau - \delta, \tau]$. Then, ϕ undergoes exactly h transitions over $t \oplus I$, for some $h \leq k \leq n$. If $b(t + \tau) \models \phi$ then, from Lemma 8, $b(t) \models \mathsf{happ}(\phi, h+1, [0, \tau])$, which implies $b(t) \models \mathsf{yieldsT}(\phi, n+1, [0, \tau])$ according to (8). Conversely, if $b(t) \models \mathsf{yieldsT}(\phi, n+1, [0, \tau])$ then $b(t) \models \mathsf{happ}(\phi, \widetilde{h}, [0, \tau])$ for some $\widetilde{h} \leq n+1$. In particular, it is $b(t) \models \mathsf{happ}(\phi, h+1, I)$; hence $b(t + \tau) \models \phi$ from Lemma 8. □

4.3 Syntactic Characterizations

This section defines non-Berkeleyness and bounded variability syntactically.

4.3.1 Non-Berkeleyness

Behaviors. The following formula χ_δ characterizes behaviors that are non-Berkeley for $\delta > 0$, that is $b \in \mathcal{B}\Sigma\mathbb{I}_\delta$ with $\Sigma = 2^{\mathcal{P}}$ iff $b \models \chi_\delta$.

$$\chi_\delta \quad = \quad \mathrm{Alw}\left(\Diamond_{[0,\delta]}\left(\bigvee_{\beta \in 2^{\mathcal{P}}} \Box_{[0,\delta]}(\beta) \right) \wedge \left(\overleftarrow{\Box}(\bot) \Rightarrow \bigvee_{\beta \in 2^{\mathcal{P}}} \Box_{[0,\delta]}(\beta) \right) \right)$$

Note that the second conjunct is needed only for time domains bounded to the left, where it holds precisely at the origin.

While χ_δ has size exponential in $|\mathcal{P}|$, it is possible to express non-Berkeleyness with a formula which is polynomial in $|\mathcal{P}|$. This will be useful when assessing the complexity class of MTL over non-Berkeley behaviors (in Section 6). To this end, let us first define:

$$\begin{aligned}
\mathsf{RT}(\beta) &= \triangle(\beta) \wedge \beta \quad \vee \quad \triangle(\neg\beta) \wedge \neg\beta \\
\mathsf{LT}(\beta) &= \triangle(\beta) \wedge \neg\beta \quad \vee \quad \triangle(\neg\beta) \wedge \beta \\
\mathsf{GT}(\beta) &= \triangle(\beta) \quad \vee \quad \triangle(\neg\beta)
\end{aligned}$$

that model a right-continuous, left-continuous, and generic transition of β, respectively. Then, we introduce:

$$\chi_\delta^{\mathrm{R}} \quad = \quad \bigwedge_{\beta \in \Sigma} \left(\mathsf{RT}(\beta) \Rightarrow \bigwedge_{\gamma \in \Sigma} \left(\begin{array}{c} \Box_{(0,\delta)}(\neg \mathsf{GT}(\gamma)) \\ \wedge \\ \mathsf{GT}(\gamma) \Rightarrow \mathsf{RT}(\gamma) \\ \wedge \\ \Diamond_{(0,\delta]}(\mathsf{GT}(\gamma)) \Rightarrow \Diamond_{(0,\delta]}(\mathsf{LT}(\gamma)) \end{array} \right) \right)$$

$$\chi_\delta^{\mathrm{L}} \quad = \quad \bigwedge_{\beta \in \Sigma} \left(\mathsf{LT}(\beta) \Rightarrow \bigwedge_{\gamma \in \Sigma} \left(\begin{array}{c} \Box_{(0,\delta]}(\neg \mathsf{GT}(\gamma)) \\ \wedge \\ \mathsf{GT}(\gamma) \Rightarrow \mathsf{LT}(\gamma) \end{array} \right) \right)$$

$$\chi_\delta^{\mathrm{I}} = \overleftarrow{\Box}(\bot) \Rightarrow \bigwedge_{\beta \in \Sigma} \Box_{[0,\delta]}(\beta \vee \neg\beta)$$

$$\chi'_\delta \quad = \quad \mathrm{Alw}\big(\chi_\delta^{\mathrm{R}} \wedge \chi_\delta^{\mathrm{L}} \wedge \chi_\delta^{\mathrm{I}}\big)$$

χ_δ^{R} describes the non-Berkeley requirement about a right-continuous transition: no other transition can occur over $(0, \delta)$, if there is a transition at the current instant it must also be right-continuous, and if there is a transition at δ it must be left-continuous, so that a closed interval of size δ is fully contained between the two consecutive transitions. Similarly, χ_δ^{R} describes the non-Berkeley requirement about a left-continuous transition. Finally, χ_δ^{I} describes the non-Berkeley requirement at the origin of a time domain bounded to the left. It should be clear that $b \in \mathcal{B}\Sigma\mathbb{I}_\delta$ with $\Sigma = 2^{\mathcal{P}}$ iff $b \models \chi'_\delta$, and χ'_δ has size quadratic in $|\mathcal{P}|$.

Words. The following formula $\chi_\delta^{\mathcal{W}}$ characterizes words that are non-Berkeley for $\delta > 0$, that is $w \in \mathcal{T}\Sigma\mathbb{I}_\delta^\omega \cup \mathcal{T}\Sigma\mathbb{I}_\delta^*$ with $\Sigma = 2^{\mathcal{P}}$ iff $w \models \chi_\delta^{\mathcal{W}}$.

$$\chi_\delta^{\mathcal{W}} \quad = \quad \mathrm{Alw}\big(\bigcirc_{\geq \delta}(\top) \vee \neg\bigcirc(\top)\big)$$

Note that the second conjunct holds iff it is evaluated at the last position in a finite word.

4.3.2 Bounded Variability

Behaviors. To describe bounded variability syntactically over behaviors, we first introduce QITL formula $\mathsf{pt}(k, I)$, for $k > 0$.

$$\mathsf{pt}(k, I) = \diamondsuit_I^k \left(\bigvee_{\beta \in \mathcal{P}} \Big(\triangle(\beta) \wedge \bigcirc(\neg\beta) \;\vee\; \triangle(\neg\beta) \wedge \bigcirc(\beta) \Big) \right) \wedge \neg \diamondsuit_I^{k+1} \left(\bigvee_{\beta \in \mathcal{P}} \Big(\triangle(\beta) \wedge \bigcirc(\neg\beta) \;\vee\; \triangle(\neg\beta) \wedge \bigcirc(\beta) \Big) \right)$$

If we let $\mathsf{pt}(0, I) = \neg \diamondsuit_I^1 \left(\bigvee_{\beta \in \mathcal{P}} \Big(\triangle(\beta) \wedge \bigcirc(\neg\beta) \vee \triangle(\neg\beta) \wedge \bigcirc(\beta) \Big) \right)$, $\mathsf{pt}(k, I)$ states that there are exactly $k \geq 0$ *punctual* transitions of atomic propositions over interval I.

Second, we introduce QITL formula $\mathsf{gt}(k, I)$, for $k > 0$:

$$\mathsf{gt}(k, I) = \diamondsuit_I^k \left(\bigvee_{\beta \in \mathcal{P}} \Big(\triangle(\beta) \vee \triangle(\neg\beta) \Big) \right) \wedge \neg \diamondsuit_I^{k+1} \left(\bigvee_{\beta \in \mathcal{P}} \Big(\triangle(\beta) \vee \triangle(\neg\beta) \Big) \right)$$

If we let $\mathsf{gt}(0, I) = \neg \diamondsuit_I^1 \left(\bigvee_{\beta \in \mathcal{P}} \Big(\triangle(\beta) \vee \triangle(\neg\beta) \Big) \right)$, $\mathsf{gt}(k, I)$ states that there are exactly $k \geq 0$ (generic, i.e., punctual or not) transitions of atomic propositions over interval I.

Finally, the following formula $\chi_{k,\delta}$ characterizes behaviors with variability bounded by k, δ, that is $b \in \mathbb{B}\Sigma\mathbb{T}_{k,\delta}$ with $\Sigma = 2^{\mathcal{P}}$ iff $b \models \chi_{k,\delta}$.

$$\chi_{k,\delta}^{\mathsf{G}} = \bigvee_{\substack{0 \leq j \leq k \\ 0 \leq h \leq \lfloor j/2 \rfloor}} \mathsf{pt}(h, [0, \delta]) \wedge \mathsf{gt}(j - h, [0, \delta])$$

$$\chi_{k,\delta}^{\mathsf{I}} = \overleftarrow{\Box}(\bot) \wedge \bigvee_{\beta \in \mathcal{P}} \begin{pmatrix} \beta \wedge \bigcirc(\neg\beta) \\ \vee \\ \neg\beta \wedge \bigcirc(\beta) \end{pmatrix} \Rightarrow \bigvee_{\substack{0 \leq j \leq k-2 \\ 0 \leq h \leq \lfloor j/2 \rfloor}} \begin{pmatrix} \mathsf{pt}(h, (0, \delta]) \\ \wedge \\ \mathsf{gt}(j - h, (0, \delta]) \end{pmatrix}$$

$$\chi_{k,\delta} = \mathrm{Alw}\big(\chi_{k,\delta}^{\mathsf{G}} \wedge \chi_{k,\delta}^{\mathsf{I}}\big)$$

More precisely, $\chi_{k,\delta}^{\mathsf{G}}$ applies to any time instant and requires that at most k transitions (weighted according to whether they are punctual or not) occur over any closed interval of size δ. On the other hand, $\chi_{k,\delta}^{\mathsf{I}}$ applies only at the origin of time domains that are bounded to the left: if there is a punctual transition at the origin, there must be at most $k - 2$ transitions over the residual interval $(0, \delta]$ (in fact, $\lim_{t \to 0^-} b(t)$ is undefined and hence different than $b(0)$); if not, it is clear that the general formula $\chi_{k,\delta}^{\mathsf{G}}$ is enough. Note that the size of $\chi_{k,\delta}$ is polynomial in $|\mathcal{P}|, k$.

5 Decidability Results

For simplicity, in this section we assume future-only MTL formulas. It is however clear that the results can be extended to MTL with past operators by providing a few additional details. We also assume formulas in normal form (introduced in Section 3.1).

5.1 MTL over Non-Berkeley Behaviors

This section shows that MTL is decidable over non-Berkeley behaviors, by providing a translation from MTL formulas to QITL formulas. The translation is introduced first for the simpler case of flat MTL formulas.

5.1.1 From Flat MTL to MITL

Every *flat* MTL formula (i.e., where no temporal operators are nested) can be translated into an MITL formula that is equivalent over non-Berkeley behaviors. In particular, given that any MITL formula is also an MTL formula, we only need to prove that the following equivalence holds over behaviors $b \in \mathcal{B}\Sigma\mathbb{T}_\delta$, where β is a Boolean combination of atomic propositions.

$$\Diamond_{=\mathrm{d}}(\beta) \equiv \begin{cases} \Diamond_{[\mathrm{d}-\delta,\mathrm{d}]}\Big(\Box_{[0,\delta]}(\beta)\Big) & \text{if } \mathrm{d} > \delta \\ \Box_{[0,\mathrm{d}]}(\beta) \vee \Diamond_{[0,\mathrm{d}]}\Big(\Box_{[0,\delta]}(\beta)\Big) & \text{if } \mathrm{d} \leq \delta \end{cases} \tag{9}$$

Proof. Let us start with the simpler $\Leftarrow$ direction, and let t be the current instant. If $\mathrm{d} > \delta$, there exists a $t' \in [t+\mathrm{d}-\delta, t+\mathrm{d}]$ such that β holds over $I_\beta = [t', t'+\delta]$. It suffices to show that $t+\mathrm{d} \in I_\beta$; in fact $t+\mathrm{d} \in I_\beta$ iff $t' \leq t+\mathrm{d} \leq t'+\delta$ iff $t+\mathrm{d} \leq t+\mathrm{d} \leq t+\mathrm{d}-\delta+\delta$, which is clearly satisfied. If $\mathrm{d} \geq \delta$ and $\Box_{[0,d]}(\beta)$, clearly β holds at $t+d$ in particular. If instead $\Diamond_{[0,\mathrm{d}]}\Big(\Box_{[0,\delta]}(\beta)\Big)$, there exists a $t' \in [t, t+\mathrm{d}]$ such that β holds over $[t', t'+\delta]$; then it is easy to check that $t+\mathrm{d} \in [t', t'+\delta]$.

Let us now consider the $\Rightarrow$ direction, and let t be the current instant. Assume first $\mathrm{d} > \delta$, so the interval $[\mathrm{d}-\delta, d]$ is non-empty (and non-punctual). Let us consider χ_δ at $t+\mathrm{d}$, where β holds: there exists $t' \in [t+\mathrm{d}-\delta, t+\mathrm{d}]$ such that β holds over $[t', t'+\delta]$. This can be expressed equivalently as $\Diamond_{[\mathrm{d}-\delta,\mathrm{d}]}\Big(\Box_{[0,\delta]}(\beta)\Big)$ with respect to the current instant t. Let us now assume $\mathrm{d} \leq \delta$ and also that $\Box_{[0,\mathrm{d}]}(\beta)$ is false at t; thus there exists a $t' \in [t, t+\mathrm{d}]$ where β is false. Since we are assuming that β holds at $t+\mathrm{d}$, β must become true somewhere between t' and $t+\mathrm{d}$. Formally, either (a) there exists a $t'' \in (t', t+\mathrm{d}]$ such that β is false over $[t', t'')$ and is true at t''; or (b) there exists a $t'' \in [t', t+\mathrm{d})$ such that β is false over $[t', t'']$ and is true over $(t'', t''+\epsilon)$ for some $\epsilon > 0$. As usual the two cases correspond to β switching from false to true right-continuously (in (a)) or left-continuously (in (b)). Let us first consider (a), and evaluate χ_δ at t''. Since β is false to the left of t'', it must be $\Box_{[0,\delta]}(\beta)$ at t''. Since $t'' \in (t', t+\mathrm{d}]$ and $t' \in [t, t+\mathrm{d}]$, we have $\Diamond_{[0,\mathrm{d}]}\Big(\Box_{[0,\delta]}(\beta)\Big)$ at t *a fortiori*. Let us now consider (b), and evaluate χ_δ "arbitrarily close to" t'', from the right to the left. This implies that $\Box_{[0,\delta]}(\beta)$ also holds "arbitrarily close to" t'', so β holds over $(t'', t''+\delta+\nu)$ for some $\nu > 0$ (as it must contain a closed interval of length δ). Since $t'' \in [t', t+\mathrm{d})$ and $t' \in [t, t+\mathrm{d}]$, we have $\Diamond_{[0,\mathrm{d}]}\Big(\Box_{[0,\delta]}(\beta)\Big)$ at t *a fortiori*. $\Box$

5.1.2 From MTL to QITL

For generic MTL formulas ϕ, (9) does not hold. However, we provide the following equivalence for distance formulas in generic MTL formulas over behaviors that are non-Berkeley for some $\delta > 0$.

Lemma 10. *For any MTL formula ϕ over any behavior $b \in \mathcal{B}_{\Sigma\mathbb{T}\delta}$, we have:*

$$\Diamond_{=\mathrm{d}}(\phi) \quad \equiv \quad \mathsf{yieldsT}(\phi, \theta(\phi) + 1, [\max(0, \mathrm{d} - \delta), \mathrm{d}])$$

Proof. Let $I = [\max(0, \mathrm{d} - \delta), \mathrm{d}]$. From Lemma 7, ϕ undergoes at most $\theta(\phi)$ transitions over $t \oplus I$. So, from Lemma 9, we have immediately that $b(t+\mathrm{d}) \models \phi$ — i.e., $b(t) \models \Diamond_{=\mathrm{d}}(\phi)$ — iff $b(t) \models \mathsf{yieldsT}(\phi, \theta(\phi) + 1, I)$. □

5.1.3 Decidability of MTL over Non-Berkeley Behaviors

It is now straightforward to prove the decidability of MTL over non-Berkeley behaviors. To this end, let us introduce the following translation function μ from MTL formulas to QITL formulas, where ψ is any MITL formula and ϕ is any MTL formula.

$$\begin{aligned}
\mu(\psi) &\equiv \psi \\
\mu(\neg\phi) &\equiv \neg\mu(\phi) \\
\mu(\phi_1 \wedge \phi_2) &\equiv \mu(\phi_1) \wedge \mu(\phi_2) \\
\mu(\mathsf{U}(\phi_1, \phi_2)) &\equiv \mathsf{U}(\mu(\phi_1), \mu(\phi_2)) \\
\mu(\Diamond_{<\mathrm{d}}(\phi)) &\equiv \Diamond_{<\mathrm{d}}(\mu(\phi)) \\
\mu(\Diamond_{=\mathrm{d}}(\phi)) &\equiv \mathsf{yieldsT}(\mu(\phi), \theta(\phi) + 1, [\max(0, \mathrm{d} - \delta), \mathrm{d}])
\end{aligned}$$

Given the "non-standard" nature of the full QITL language, it may be useful to introduce its operators in the translation of MTL formulas only when strictly needed. To this end, we provide an alternative translation ν as follows.

$$\begin{aligned}
\nu(\psi) &\equiv \psi \\
\nu(\Diamond_{=\mathrm{d}}(\beta)) &\equiv \begin{cases} \Diamond_{[\mathrm{d}-\delta,\mathrm{d}]}\left(\Box_{[0,\delta]}(\beta)\right) & \text{if } \mathrm{d} > \delta \\ \Box_{[0,\mathrm{d}]}(\beta) \vee \Diamond_{[0,\mathrm{d}]}\left(\Box_{[0,\delta]}(\beta)\right) & \text{if } \mathrm{d} \leq \delta \end{cases} \\
\nu(\Diamond_{=\mathrm{d}}(\mathsf{U}(\phi_1, \phi_2))) &\equiv \mathsf{yieldsT}(\mathsf{U}(\nu(\phi_1), \nu(\phi_2)), \theta(\mathsf{U}(\phi_1, \phi_2)) + 1, [\max(0, \mathrm{d} - \delta), \mathrm{d}]) \\
\nu(\Diamond_{=\mathrm{d}}\left(\Diamond_{\langle a,b\rangle}(\phi)\right)) &\equiv \Diamond_{\langle \mathrm{d}+a, \mathrm{d}+b\rangle}(\nu(\phi)) \\
\nu(\Diamond_{=\mathrm{d}_1}(\Diamond_{=\mathrm{d}_2}(\phi))) &\equiv \nu(\Diamond_{=\mathrm{d}_1+\mathrm{d}_2}(\nu(\phi))) \\
\nu(\Diamond_{=\mathrm{d}}(\neg\phi)) &\equiv \neg\nu(\Diamond_{=\mathrm{d}}(\phi)) \\
\nu(\Diamond_{=\mathrm{d}}(\phi_1 \wedge \phi_2)) &\equiv \nu(\Diamond_{=\mathrm{d}}(\phi_1)) \wedge \nu(\Diamond_{=\mathrm{d}}(\phi_2))
\end{aligned}$$

For both translations, it is straightforward to prove the following equivalence result.

Theorem 11. *For any MTL formula ϕ, for any behavior $b \in \mathcal{B}_{\Sigma\mathbb{T}_\delta}$ for some $\delta > 0$, the QITL formulas $\mu(\phi), \nu(\phi)$ are such that: $b \models_{\mathbb{T}} \phi$ iff $b \models_{\mathbb{T}} \mu(\phi)$ iff $b \models_{\mathbb{T}} \nu(\phi)$.*

Proof. The proof is trivial by induction on the structure of ϕ, from (9) and Lemma 10. □

Theorem 11, the decidability of MITL and QITL [AFH96, HR04], and the syntactic characterization of non-Berkeleyness by means of the χ_δ formula, immediately imply the following.

Corollary 12. *For any $\delta > 0$, the satisfiability of MTL formulas is decidable over $\mathcal{B}\Sigma\mathbb{T}_\delta$.*

Proof. Given a generic MTL formula ϕ, ϕ is satisfiable over $\mathcal{B}\Sigma\mathbb{T}_\delta$ iff $\phi' = \phi \wedge \chi'_\delta$ is satisfiable over non-Zeno behaviors. In turn, by Theorem 11, ϕ' is satisfiable over non-Zeno behaviors iff $\phi'' = \mu(\phi) \wedge \chi'_\delta$ is. Since ϕ'' is a QITL formula, the theorem follows from Proposition 4. □

5.2 MTL over Bounded Variably Behaviors

The results of the previous section can be extended to the case of behaviors with bounded variability along the following lines. First, consider the claim: for any $b \in \mathcal{B}\Sigma\mathbb{T}_{k,\delta}$ and MTL formula ϕ, it is $b_\phi \in \mathcal{B}\Sigma\mathbb{T}_{k+\theta(\phi),\delta}$. The claim can be proved similarly as for Lemma 7, where the base case for Boolean combinations β is changed into $2 + k$, whereas the inductive steps are essentially unaffected, provided the inductive hypothesis about the variability being bounded by θ is replaced by it being bounded by $\theta + k$. Correspondingly, we can introduce a translation μ' from MTL to QITL formulas which is obtained from μ by replacing $\theta(\phi)$ with $k + \theta(\phi)$. Finally, QITL formula $\mu'(\phi) \wedge \chi_{k,\delta}$ is satisfiable over $\mathcal{B}\Sigma\mathbb{T}$ iff ϕ is satisfiable over $\mathcal{B}\Sigma\mathbb{T}_{k,\delta}$. Hence, MTL is decidable over $\mathcal{B}\Sigma\mathbb{T}_{k,\delta}$.

5.3 MTL over Words

The decidability of MTL over words with bounded variability has been already been proved by Wilke [Wil94]. More precisely, Wilke's results entail the satisfiability of MTL over $\mathcal{T}\Sigma\mathbb{T}^\omega_{k,1}$. However, it is clear that any MTL formula ϕ is satisfiable over $\mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta}$ with $\delta = n/d$ iff ϕ' is satisfiable over $\mathcal{T}\Sigma\mathbb{T}^\omega_{dk,1}$ where ϕ' is obtained from ϕ by scaling all its constants by n. From the equivalences $\mathcal{T}\Sigma\mathbb{T}^\omega_\delta = \mathcal{T}\Sigma\mathbb{T}^\omega_{1,\delta}$ and $\mathcal{T}\Sigma\mathbb{T}^*_\delta = \mathcal{T}\Sigma\mathbb{T}^*_{1,\delta}$ (Proposition 1) it is clear that MTL is decidable over non-Berkeley words as well.

Additionally, we can prove the same result for non-Berkeley words with a different technique. Namely, we provide a translation from MTL to MITL that preserves satisfiability for non-Berkeley words. Shortly, we have the following equivalence over non-Berkeley words for MTL formulas:

$$\Diamond_{=\mathrm{d}}(\phi) \equiv \begin{cases} \Box_{(\mathrm{d}-\delta,\delta)}(\bot) \wedge \Diamond_{(\mathrm{d}-\delta,\delta]}(\phi) & \text{if } \mathrm{d} \geq \delta \\ \bot & \text{if } \mathrm{d} < \delta \end{cases} \tag{10}$$

Then, if we replace every occurrence of the $\Diamond_{=\mathrm{d}}$ operator in an MTL formula according to (10) we obtain an MITL formula which is equivalent to the original MTL formula over words that are non-Berkeley for $\delta > 0$. The syntactic capability of expressing non-Berkeleyness requirements in MITL entails the decidability of MTL over $\mathcal{T}\Sigma\mathbb{T}^\omega_\delta \cup \mathcal{T}\Sigma\mathbb{T}^*_\delta$.

Extending this approach to $\mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta} \cup \mathcal{T}\Sigma\mathbb{T}^*_{k,\delta}$ seems to require some kind of "counting" modality as in the case of behaviors. Wilke [Wil94] showed that words in $\mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta} \cup \mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta}$ can be characterized by means of a formula in a monadic second-order logic that is decidable and expressively complete for Alur and Dill's timed automata [AD94]. This entails that it is possible to introduce suitable "counting" modalities that do not compromise decidability. We leave the exploration of this possibility to future work.

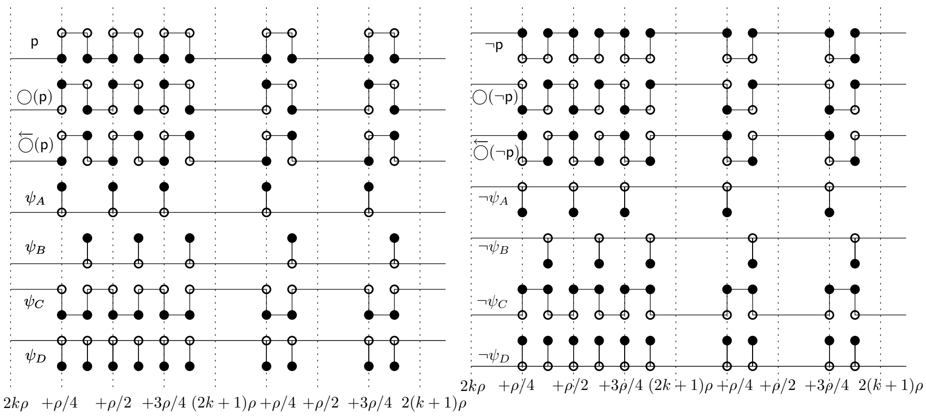

Figure 3: All behaviors of MITL over $b^{\rho,\epsilon}$ for $\epsilon = \rho/8$.

6 Related Results

This section discusses the expressiveness and complexity of MTL over non-Berkeley and bounded variably behaviors and words.

6.1 Expressiveness of MTL over non-Berkeley

The technique used in Section 5 to assess the decidability of MTL over non-Berkeley behaviors involved the translation of MTL formulas into QITL, a strict superset of MITL. This raises the obvious question of whether QITL is really needed in translating MTL to a decidable logic. We provide a partial negative answer to this question, by showing that MITL is strictly less expressive than MTL over non-Berkeley behaviors,[5] and consequently we cannot generalize MITL formula (9) to handle generic MTL formulas. This answer is only partial because we address expressiveness, not equi-satisfiability; that is, it might be possible to construct, for every MTL formula, a corresponding MITL formula which is equi-satisfiable over non-Berkeley behaviors but requires additional atomic propositions to be built.

Let us consider explicitly full MITL (and not just future-only MITL) because MITL increases its expressive power if we add past operators [MNP05, AH92]. Expressiveness separation results are usually quite convoluted. The ensuing complexity is commonly tamed by considering behaviors with pointwise transitions only, as it is usually simpler to characterize exhaustively the truth value of formulas over such behaviors. This is however impossible for non-Berkeley behaviors (where pointwise transitions cannot occur), so the following separation proofs would be quite verbose if all details were spelled out. In fact, consider behavior $b^{\rho,\epsilon}$ defined as $\mathsf{p} \in b^{\rho,\epsilon}(t)$ iff $t \in \bigcup_{k\in\mathbb{Z}} \big(\rho(k+1/4) \oplus (0,\epsilon) \cup \rho(k+3/4) \oplus (0,\epsilon) \cup \rho(2k+1/2) \oplus (0,\epsilon)\big)$; we have the following.

Lemma 13. *Let $0 < \epsilon < \rho/4$; then the the truth value over $b^{\rho,\epsilon}$ of any MITL formula ϕ of granularity ρ coincides with one of those in $\mathcal{G} \cup \neg\mathcal{G}$, where $\mathcal{G} =$*

[5] The difference in expressive power between MTL and MITL is obvious over generic behaviors, where MITL is decidable while MTL is not.

$\{\top, \mathsf{p}, \bigcirc(\mathsf{p}), \overleftarrow{\bigcirc}(\mathsf{p}), \psi_A = \neg\mathsf{p} \wedge \bigcirc(\mathsf{p}), \psi_B = \overleftarrow{\bigcirc}(\mathsf{p}) \wedge \bigcirc(\neg\mathsf{p}), \psi_C = \overleftarrow{\bigcirc}(\neg\mathsf{p}) \wedge \bigcirc(\neg\mathsf{p}), \psi_D = \neg\psi_A \wedge \neg\psi_B\}$ *and* $\mathcal{G}' = \{\neg\psi | \psi \in \mathcal{G}\}$ *(see Figure 3).*

Proof. The proof is straightforward — albeit quite tedious — by induction on the structure of ϕ.

Consider for instance $\phi = \mathsf{U}(\gamma_1, \gamma_2)$ with $\gamma_1 \equiv \mathsf{p}$ and $\gamma_2 \equiv \neg\mathsf{p}$; ϕ holds precisely when p is true and will become false right-continuously. However, this is the case whenever p holds, as well as when p becomes true left-continuously. Therefore ϕ is equivalent to $\bigcirc(\mathsf{p})$.

Another example is for $\phi = \Diamond_{<\mathrm{d}}(\gamma)$ with $\mathrm{d} = k\rho$ for $k \geq 1$ and $\gamma = \bigcirc(\mathsf{p})$. It is simple to check that any interval of size $\geq \rho$ encompasses points where $\bigcirc(\mathsf{p})$ holds. Correspondingly, ϕ is simply equivalent to $\top$.

The other cases are handled similarly. □

Correspondingly, we have a first partial separation result.

Lemma 14. *For any* $\delta > 0$*, for all* $\rho > 8\delta$*,* $\mathrm{MTL}^{\geq\rho}$ *is strictly more expressive than* $\mathrm{MITL}^{\geq\rho}$ *over* $\mathbb{BEH}_\delta$*.*

Proof. To prove the lemma, we fix a granularity $\rho > 8\delta$ and show that there exists an MTL^ρ formula $\phi^\rho_{\mathrm{Dist}}$ which is equivalent to no $\mathrm{MITL}^{\geq\rho}$ formula over some behavior $b^{\rho,\epsilon} \in \mathbb{BEH}_\delta$. Consider MTL formula $\mathrm{MTL}^\rho \ni \phi^\rho_{\mathrm{Dist}} = \mathsf{p} \Rightarrow \Diamond_{=\rho}(\mathsf{p})$ whose granularity is ρ. You can check that $b^{\rho,\epsilon}(t) \not\models \phi^\rho_{\mathrm{Dist}}$ iff $t \in \rho(2k + 1/2) \oplus (0, \epsilon)$ for some $k \in \mathbb{Z}$. However, Lemma 13 showed that no $\mathrm{MITL}^{\geq\rho}$ formula has this behavior. □

Lemma 14 settles the problem of expressiveness only for formulas of granularity ρ with respect to behaviors that are non-Berkeley for $\delta < \rho/8$. On the contrary, we are interested also in determining if the same relations holds when the behaviors are "slower than the granularity", that is for $\rho \leq 8\delta$ and in particular if $\rho < \delta$. In this case, a full characterization of MITL formulas is even more tedious than the one in Lemma 13. On the one hand, in order to consider behaviors that are non-trivial in that punctuality provides indeed more expressiveness, one has to take behavior with a "long" period. On the other hand, such "slower" behaviors give rise to many different "derived" behaviors, where by derived we mean those representing the truth of some MITL formula. We claim that MTL is nonetheless more expressive than MITL even in such cases. For simplicity, we only sketch a proof idea for this case, and leave all details for a longer version of the paper.

Lemma 15. *For any* $\delta > 0$*, for all* $\rho < \delta$*,* $\mathrm{MTL}^{\geq\rho}$ *is strictly more expressive than* $\mathrm{MITL}^{\geq\rho}$ *over* $\mathbb{BEH}_\delta$*.*

Proof sketch. The idea is to build a behavior c^ρ and a formula ψ^ρ_b such that the truth value $b_{\psi^\rho_b}$ of ψ^ρ_b over c^ρ is qualitatively similar to the $b^{\rho,\epsilon}$ of Lemma 13. Then, we build a punctual formula which has ψ^ρ_b as a subformula. Similarly to what happens in Lemma 14 for $\phi^\rho_{\mathrm{Dist}}$, the new formula has a behavior different than any $\mathrm{MITL}^{\geq\rho}$ formula over c^ρ. Hence, we have the corresponding expressiveness separation.

More precisely, consider behavior c^ρ defined as $\mathsf{p} \in c^\rho(t)$ iff $t \in \bigcup_{\substack{i \in \mathbb{Z} \\ j \in [0..3]}} 9\rho(i + (2j+1)/8) \oplus [0, 9\rho/8)$. It is pictured in Figure 4(a), over a period of 9ρ units.

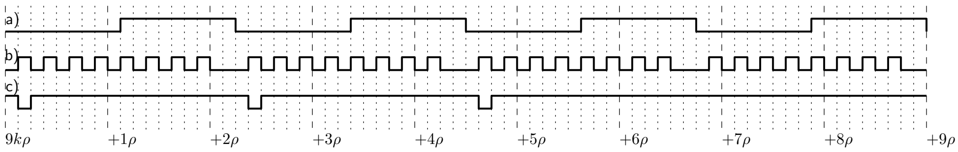

Figure 4: a) c^ρ; b) $b_{\psi_b^\rho}$; c) $b_{\psi_b^\rho \Rightarrow \Diamond_{=2\rho}(\psi_b^\rho)}$.

Let ψ_b^ρ be MITL formula $\bigvee_{0 \leq k \leq 7} \Box_{k\rho \oplus [0,1)}(\mathsf{p})$. The truth behavior $b_{\psi_b^\rho}$ of ψ_b^ρ is picture in Figure 4(b). Now, let us consider MTL$^\rho$ formula $\psi_b^\rho \Rightarrow \Diamond_{=2\rho}(\psi_b^\rho)$; its truth value changes over time as in Figure 4(c). Notice that we have introduced a sort of "irregular" behavior as with ϕ_{Dist}^ρ in Lemma 13. Then, a tedious case analysis of all MITL$^{\geq\rho}$ formulas would show that none of them has the same behavior as $\psi_b^\rho \Rightarrow \Diamond_{=2\rho}(\psi_b^\rho)$ over c^ρ, which entails the theorem. ☐

The desired separation result follows as a corollary of the previous lemma.

Theorem 16. *For any $\delta > 0$, MTL is strictly more expressive than MITL over $\mathbb{B}\Sigma\mathbb{T}_\delta$.*

Proof. Note that MITL$^{\rho_1} \supseteq$ MITL$^{\rho_2}$ for all $\rho_1 = \rho_2/k$ with k any positive integer. Hence, MITL $= \bigcup_{k \in \mathbb{N}_{>0}}$ MITL$^{\geq \rho/k}$ for any $\rho > 0$. Then, the theorem follows from Lemma 15. ☐

6.2 Complexity of MTL over Non-Berkeley

This section shows that the satisfiability problem for MTL formulas over non-Berkeley (and bounded variably) behaviors has the same complexity as the same problem for MITL over generic behaviors.

Theorem 17. *The satisfiability problem for MTL over behaviors in $\mathbb{B}\Sigma\mathbb{T}_\delta$ is* **EXPSPACE***-complete (assuming a succinct encoding of constants for MTL formulas).*

Proof. The fact that the problem is in **EXPSPACE** follows from the translation procedure of Section 3 from an MTL formula ϕ to an equi-satisfiable QITL formula of size polynomial in $|\phi|$, and from the complexity of QITL (Proposition 4).

The **EXPSPACE**-hardness of MTL satisfiability over non-Berkeley behaviors can be proved by reducing the corresponding problem over the integers. Let ϕ be any MTL formula. It is always possible to build an equi-satisfiable formula $\overline{\phi}$ obtained from ϕ by "flattening" nesting temporal operators through the introduction of additional fresh atomic propositions (see [FS08] for details of the straightforward construction). Hence, $\overline{\phi}$ can be defined by $\phi ::= \beta \mid \alpha \vee \mathsf{U}_J(\beta_1, \beta_2) \mid \alpha \vee \mathsf{R}_J(\beta_1, \beta_2) \mid \alpha \vee \Diamond_{=\mathrm{d}}(\beta) \mid \phi_1 \vee \phi_2 \mid \phi_1 \wedge \phi_2$ with α an atomic proposition (negated or unnegated), β a Boolean combination of atomic propositions, and J a nonsingular interval. Correspondingly, every occurrence of an atomic proposition α in $\overline{\phi}$ is called *existential* iff one of the following is the case: (1) the occurrence appears in the second argument β_2 of a subformula in the form $\mathsf{U}_J(\beta_1, \beta_2)$; (2) the occurrence appears in the first argument β_1 of a subformula in the form $\mathsf{R}_J(\beta_1, \beta_2)$. Every occurrence that is not existential is called *universal*.

We reduce the satisfiability of $\overline{\phi}$ (and hence of ϕ) over the integers to the satisfiability of another formula ξ over non-Berkeley behaviors for some $\delta > 0$. To this end, let $\theta = k\delta$ be a multiple of δ for some $k \geq 2$. We introduce a fresh proposition clock; its behavior is defined by the formula:

$$\begin{aligned}\pi \quad = \quad & \Box_{(0,\theta]}(\mathsf{clock}) \wedge \mathrm{Alw}\Big(\Box_{(0,\theta]}(\mathsf{clock}) \Leftrightarrow \Box_{(\theta,2\theta]}(\neg\mathsf{clock})\Big) \\ & \wedge \mathrm{Alw}\Big(\Box_{(0,\theta]}(\neg\mathsf{clock}) \Leftrightarrow \Box_{(\theta,2\theta]}(\mathsf{clock})\Big)\end{aligned}$$

which describes clock as a square wave with a 50% duty cycle and where transitions are left-continuous.

Then, let ϕ' be the MTL formula obtained from $\overline{\phi}$ by substituting every existential occurrence of an atomic proposition $\alpha = \mathsf{p}$ or $\alpha = \neg\mathsf{p}$ by $(\neg\mathsf{clock} \wedge \bigcirc(\mathsf{clock})) \wedge \bigcirc(\alpha)$ and every universal occurrence by $(\neg\mathsf{clock} \wedge \bigcirc(\mathsf{clock})) \Rightarrow \bigcirc(\alpha)$, and by multiplying all constants in $\overline{\phi}$ by 2θ. Finally, let $\xi = \pi \wedge \phi'$. We claim that ξ is satisfiable over non-Berkeley behaviors for δ iff ϕ is satisfiable over the integers. Let us first consider a behavior i over the integers such that $i \models_{\mathbb{Z}} \phi$ and thus also $i \models_{\mathbb{Z}} \overline{\phi}$ by construction. Consider behavior r over the reals built as follows. First, $\mathsf{clock} \in r(t)$ iff $t \in (2k\theta, (2k+1)\theta]$ for some $k \in \mathbb{Z}$. Then, for any other proposition p, if $\mathsf{p} \in i(k)$ for some integer k, let p hold over $(2k\theta, (2k+1)\theta]$ over r, and let p be false over r otherwise. Now, the transformation from $\overline{\phi}$ to ϕ' is such that the truth of ϕ' depends only on what happens at raising edges of clock. Correspondingly, it is not difficult to check by induction on the structure of ϕ' that $r \models_{\mathbb{R}} \phi'$ and thus also $r \models_{\mathbb{R}} \xi$. For the converse, let r be a non-Berkeley behavior such that $r \models_{\mathbb{R}} \xi$, and thus $r \models_{\mathbb{R}} \phi'$ in particular. We build a behavior i over the integers as follows. For any proposition p and instant of time k, $\mathsf{p} \in i(k)$ iff $r(2k\theta) \models_{\mathbb{R}} \bigcirc(\mathsf{p})$; intuitively we discard whatever happens between integer multiples of 2θ. Note that, for any non-Zeno behavior b (and thus for non-Berkeley behaviors *a fortiori*), $b(t) \models_{\mathbb{R}} \neg\bigcirc(\mathsf{p})$ iff $b(t) \models_{\mathbb{R}} \bigcirc(\neg\mathsf{p})$. Consequently, thanks to how ϕ' has been built from $\overline{\phi}$ it is straightforward to show by induction that $i \models_{\mathbb{Z}} \overline{\phi}$ and thus also $i \models_{\mathbb{Z}} \phi$ by construction. □

With a very similar justification we can prove the following.

Theorem 18. *The satisfiability problem for MTL over behaviors in $\mathcal{B}\Sigma\Gamma_{k,\delta}$ is* **EXPSPACE**-*complete (assuming a unary encoding of k).*

The following is instead derivable from the previous theorem and [Wil94].

Theorem 19. *The satisfiability problem for MTL over words in $\mathcal{T}\Sigma\Gamma^{\omega}_{k,\delta} \cup \mathcal{T}\Sigma\Gamma^{*}_{k,\delta}$ for $k \geq 1$ is* **EXPSPACE**-*complete (assuming a unary encoding of k).*

Proof. The **EXPSPACE**-hardness proof can be worked out as for Theorem 17, with trivial modifications. Membership in **EXPSPACE** for $k = 1$ can also be derived as in Theorem 17, through the results of Section 5.3. Membership in **EXPSPACE** for $k > 1$ can instead be derived by examining Wilke's proof of the decidability of MTL over $\mathcal{T}\Sigma\Gamma^{\omega}_{k,\delta}$ [Wil94], where it is clear that the underlying decision procedure has the same complexity as other logics of similar expressive power (MITL, in particular). It is also clear that the same results can be extended to the case of finite words. □

7 Undecidability Results

MTL is no more decidable if we consider all non-Berkeley behaviors for any δ together. More precisely, the satisfiability problem for MTL over $\mathbb{BET}_{\exists\delta}$ is Σ^0_1-complete; compare against the same problem over $\mathbb{BET}$ where it is Σ^1_1-complete.

Theorem 20. *The satisfiability problem for MTL over behaviors in $\mathbb{BET}_{\exists\delta}$ is $\Sigma^0_1 = \mathbf{RE}$-complete.*

Proof. Let ϕ be a generic MTL formula. ϕ is satisfiable over $\mathbb{BET}_{\exists\delta}$ iff there exists a $\overline{\delta} > 0$ such that ϕ is satisfiable over $\mathbb{BET}_{\overline{\delta}}$. Given that $\mathbb{BET}_\gamma \supset \mathbb{BET}_{\overline{\delta}}$ for all $\gamma < \overline{\delta}$ (Proposition 1), and that the satisfiability of ϕ is decidable over $\mathbb{BET}_\gamma$ for any fixed $\gamma > 0$ (Corollary 12), the following procedure halts iff ϕ is satisfiable over $\mathbb{BET}_{\exists\delta}$: (1) let $d \leftarrow 1$; (2) decide if ϕ is satisfiable over $\mathbb{BET}_d$; (3) if not, let $d \leftarrow d/2$ and goto (2). This proves that the satisfiability problem for MTL over $\mathbb{BET}_{\exists\delta}$ is in $\mathbf{RE}$.

To show $\mathbf{RE}$-hardness, we reduce the halting problem for 2-counter machines to MTL satisfiability over $\mathbb{BET}_{\exists\delta}$. The key insight is that a halting computation is one where only a finite portion of the tape is used. Correspondingly it can be represented by a behavior where only a finite number of transition points lie within a finite amount of time; such behaviors are necessarily in $\mathbb{BET}_{\exists\delta}$ because the infimum over distances between transitions coincides with the minimum.

Then, we use standard techniques such as those in [AH93] with some simple modifications. A *2-counter machine* M consists of a finite control unit whose locations are in a finite set L and two unbounded variables C, D ranging over the naturals, called counters. The behavior of M is determined by the sequence of instructions corresponding to L. Every instruction belongs to one of three types: branching to a certain location upon a specific counter having value 0; incrementing a counter; and decrementing a counter (which accomplishes something only if the counter has a positive value). When a given location $l_H \in L$ is reached, the machine halts. The halting problem for M consists in determining whether M eventually halts. It is well-known that the halting problem for 2-counter machines is an undecidable, $\mathbf{RE}$-hard problem [Min67]. We reduce deciding whether a generic 2-counter machine M halts to deciding if a suitable MTL formula ϕ^M is satisfiable over $\mathbb{BET}_{\exists\delta}$.

We introduce the following propositions: s, $\{\mathsf{l}_i \mid l_i \in L\}$ for every location, and c, d to "count" values of the two counters as we will explain shortly. Configurations of M are triples $\langle l, c, d\rangle$ with $l \in L$ the current location, and $c, d \geq 0$ the current values of the two counters. Every such configuration is encoded by the value of the propositions over an interval of time $[k, k+1)$ for some $k \in \mathbb{N}$, with l holding at the beginning of the interval, and c (resp. d) transitioning to true exactly c (resp. d) times within the interval. Consecutive configurations are encoded in adjacent intervals.

More precisely, s is used to determine where each interval begins. In fact, its value is determined by the following formula, where $0 < \epsilon < 1/3$ is some chosen constant:

$$\psi^{\mathsf{s}} \quad \equiv \quad \Box_{[0,\epsilon)}(\mathsf{s}) \wedge \Box_{[\epsilon,1]}(\neg\mathsf{s}) \wedge \Box_{(1,1+\epsilon)}(\mathsf{s}) \wedge \Box_{>0}(\mathsf{s} \Leftrightarrow \Diamond_{=1}(\mathsf{s}))$$

which entails that the formula $\nabla(\mathsf{s}) \equiv \neg\mathsf{s} \wedge \bigcirc(\mathsf{s})$ holds precisely at all $k \in \mathbb{N}_{\geq 1}$.

For a proposition l_i we define $\mathsf{l}_i^\flat \equiv \mathsf{l}_i \wedge \bigwedge_{j\neq i} \neg \mathsf{l}_j$. Then, the initial configuration $\langle l_0, 0, 0\rangle$ is encoded by the following formula:

$$\psi^{\text{start}} \equiv \Box_{<\epsilon}\left(\mathsf{l}_0^\flat\right) \wedge \Box_{<1}(\neg\mathsf{c} \wedge \neg\mathsf{d})$$

Every instruction is encoded by suitably "copying" the current configuration to the next state and then possibly changing the value of the counters. In particular, an increment can always be accommodated thanks to the density of the time domain. For instance, consider an instruction l_k that increments counter c (and then moves to the next instruction in l_{k+1}. This would be encoded by the following formula, where $\nabla_{=1}(\psi) \equiv \Diamond_{=1}(\nabla(\psi))$:

$$\text{Alw}\left(\nabla(\mathsf{s}) \wedge \mathsf{l}_k^\flat \Rightarrow \left(\begin{array}{l} \nabla_{=1}\left(\mathsf{l}_{k+1}^\flat\right) \\ \wedge\ \Box_{(0,1)}(\nabla(\mathsf{c}) \Rightarrow \nabla_{=1}(\mathsf{c})) \\ \wedge\ \mathsf{U}_{(0,1)}\left(\begin{array}{l} \nabla_{=1}(\mathsf{c}) \Rightarrow \nabla(\mathsf{c})\,, \\ \neg\nabla(\mathsf{c}) \wedge \nabla_{=1}(\mathsf{c}) \\ \wedge\ \mathsf{U}(\nabla_{=1}(\mathsf{c}) \Rightarrow \nabla(\mathsf{c})\,, \nabla(\mathsf{s})) \end{array}\right) \\ \wedge\ \Box_{(0,1)}(\nabla(\mathsf{d}) \Leftrightarrow \nabla_{=1}(\mathsf{d})) \end{array}\right)\right)$$

In the consequent, the first conjunct ensures that the next location is l_{k+1}; the second conjunct asserts that in the next interval contains at least as many c transitions as the current interval; the third conjunct, combined with the second, ensures that there is exactly one more c transition in the next interval than in the current one; the fourth conjunct requires that there is exactly the same number of d transitions in the current interval and in the next one.

Finally, halting can be expressed by the formula:

$$\Diamond\left(\nabla\left(\mathsf{l}_H^\flat\right)\right)$$

Now, if M has a computation that halts after N steps, a behavior b can be built such that it models the N steps according to what required by ϕ^M. In particular, after time N all propositions can take a constant value, except for s that switches at most every ϵ time units. Let $\tau(b) = \tau_1, \tau_2, \ldots$ be the sequence of transition points of b and let Z be such that $\tau_Z \leq N < \tau_{Z+1}$; then, b is in $\mathcal{BEH}_{\min(\delta,\epsilon)/2}$ for $\delta = \min_{1\leq k\leq Z-1}(\tau_{k+1} - \tau_k)$. Conversely, it is clear that a computation that halts after N steps corresponds to any behavior $b \in \mathcal{BEH}_{\exists\delta}$ satisfying ϕ^M, where N is the least instant such that $\nabla\left(\mathsf{l}_H^\flat\right)$ holds. □

As usual, the above proof can be adapted with simple modifications to work for infinite timed words. Hence, we have the following.

Theorem 21. *The satisfiability problem for MTL over words in $\mathcal{TEH}^\omega_{\exists\delta}$ is $\Sigma^0_1 =$ **RE**-complete.*

In addition, the decidability of MTL over the classes $\mathcal{BEH}_{k,\delta}$ and $\mathcal{TEH}^\omega_{k,\delta}$ entails the following.

Corollary 22. *The satisfiability problem for MTL over $\mathcal{BEH}_{\exists k\exists\delta}$ and $\mathcal{TEH}^\omega_{\exists k\exists\delta}$ is $\Sigma^0_1 =$ **RE**-complete.*

	DECIDABILITY	COMPLEXITY
$\mathcal{B}\Sigma\mathbb{T}_\delta$, $\mathcal{B}\Sigma\mathbb{T}_{k,\delta}$	Yes	**EXPSPACE**-C
$\mathcal{B}\Sigma\mathbb{T}_{\exists\delta}, \mathcal{B}\Sigma\mathbb{T}_{\exists k \exists\delta}$	No	Σ_1^0-C
$\mathcal{B}\Sigma\mathbb{T}$	No	Σ_1^1-C
$\mathcal{T}\Sigma\mathbb{T}^\omega_\delta$, $\mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta}$	Yes	**EXPSPACE**-C
$\mathcal{T}\Sigma\mathbb{T}^\omega_{\exists\delta}, \mathcal{T}\Sigma\mathbb{T}^\omega_{\exists k\exists\delta}$	No	Σ_1^0-C
$\mathcal{T}\Sigma\mathbb{T}^\omega$	No	Σ_1^1-C
$\mathcal{T}\Sigma\mathbb{T}^*_\delta$, $\mathcal{T}\Sigma\mathbb{T}^*_{k,\delta}$	Yes	**EXPSPACE**-C
$\mathcal{T}\Sigma\mathbb{T}^*_{\exists\delta}, \mathcal{T}\Sigma\mathbb{T}^*_{\exists k\exists\delta}$	Yes	non-**PR**
$\mathcal{T}\Sigma\mathbb{T}^*$	Yes	non-**PR**

Table 1: Summary of the known results.

Proof. The **RE**-hardness proof works as in the case of Theorems 20–21, from the inclusions $\mathcal{B}\Sigma\mathbb{T}_{\exists\delta} \subset \mathcal{B}\Sigma\mathbb{T}_{\exists k\exists\delta}$ and $\mathcal{T}\Sigma\mathbb{T}^\omega_{\exists\delta} \subset \mathcal{T}\Sigma\mathbb{T}^\omega_{\exists k\exists\delta}$ in Proposition 1. Membership in **RE** follows from the decidability of MTL over the classes $\mathcal{B}\Sigma\mathbb{T}_{k,\delta}$ and $\mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta}$, with the procedure: (1) let $h \leftarrow 1$, $d \leftarrow 1$; (2) decide if ϕ is satisfiable over $\mathcal{B}\Sigma\mathbb{T}_{h,d}$ (or $\mathcal{T}\Sigma\mathbb{T}^\omega_{h,d}$); (3) if not, let $h \leftarrow 2h$, $d \leftarrow d/2$ and goto (2). □

Undecidability does not carry over to finite words, where the problem is known to be decidable [OW05, OW07].

8 Summary

Table 1 summarizes the results on the expressiveness of MTL over various semantic classes. Cells without shade host previously known results; cells with a light shade are corollaries of known results; cells with a dark shade correspond to the main results discussed and proved in this paper.

As future work, it will be interesting to investigate the practical impact of the new decidability results of this paper. This will encompass, on the one hand, experimenting with implementations of decision algorithms to evaluate their performances on practical verification problems and, on the other hand, assessing which classes of systems can be naturally described with bounded variably models.

References

[ACM02] Eugene Asarin, Paul Caspi, and Oded Maler. Timed regular expressions. *Journal of the ACM*, 49(2):172–206, 2002.

[AD94] Rajeev Alur and David L. Dill. A theory of timed automata. *Theoretical Computer Science*, 126(2):183–235, 1994.

[AFH96] Rajeev Alur, Tomás Feder, and Thomas A. Henzinger. The benefits of relaxing punctuality. *Journal of the ACM*, 43(1):116–146, 1996.

[AH92] Rajeev Alur and Thomas A. Henzinger. Back to the future: Towards a theory of timed regular languages. In *Proceedings of the 33rd Annual Symposium on Foundations of Computer Science (FOCS'92)*, pages 177–186. IEEE Computer Society Press, 1992.

[AH93] Rajeev Alur and Thomas A. Henzinger. Real-time logics: Complexity and expressiveness. *Information and Computation*, 104(1):35–77, 1993.

[BMOW07] Patricia Bouyer, Nicolas Markey, Joël Ouaknine, and James Worrell. The cost of punctuality. In *Proceedings of the 22nd IEEE Symposium on Logic in Computer Science (LICS'07)*. IEEE Computer Society, 2007.

[DP07] Deepak D'Souza and Pavithra Prabhakar. On the expressiveness of MTL in the pointwise and continuous semantics. *Formal Methods Letters (Software Tools for Technology Transfer)*, 9(1):1–4, 2007.

[FPR08] Carlo A. Furia, Matteo Pradella, and Matteo Rossi. Automated verification of dense-time MTL specifications via discrete-time approximation. In Jorge Cuéllar and Tom Maibaum, editors, *Proceedings of the 15th International Symposium on Formal Methods (FM'08)*, volume 5014 of *Lecture Notes in Computer Science*, pages 132–147. Springer-Verlag, May 2008.

[FR06] Carlo A. Furia and Matteo Rossi. Integrating discrete- and continuous-time metric temporal logics through sampling. In Eugene Asarin and Patricia Bouyer, editors, *Proceedings of the 4th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'06)*, volume 4202 of *Lecture Notes in Computer Science*, pages 215–229. Springer-Verlag, September 2006.

[Frä96] Martin Fränzle. Decidability of duration calculi on restricted model classes. Technical Report Kiel MF 21/1, Institut für Informatik und Praktische Mathematik, Christian-Albrechts Universität Kiel, July 1996.

[FS08] Carlo A. Furia and Paola Spoletini. MTL satisfiability over the integers. Technical Report 2008.2, Dipartimento di Elettronica e Informazione, Politecnico di Milano, February 2008.

[GHR94] Dov M. Gabbay, Ian Hodkinson, and Mark Reynolds. *Temporal Logic (vol. 1): mathematical foundations and computational aspects*, volume 28 of *Oxford Logic Guides*. Oxford University Press, 1994.

[HR04] Yoram Hirshfeld and Alexander Moshe Rabinovich. Logics for real time: Decidability and complexity. *Fundamenta Informaticae*, 62(1):1–28, 2004.

[Koy90] Ron Koymans. Specifying real-time properties with metric temporal logic. *Real-Time Systems*, 2(4):255–299, 1990.

[LWW07] Carsten Lutz, Dirk Walther, and Frank Wolter. Quantitative temporal logics over the reals: PSPACE and below. *Information and Computation*, 205(1):99–123, 2007.

[Min67] Marvin Lee Minsky. *Computation: Finite and Infinite Machines*. Prentice Hall, 1967.

[MNP05] Oded Maler, Dejan Nickovic, and Amir Pnueli. Real time temporal logic: Past, present, future. In Paul Petterson and Wang Yi, editors, *Proceedings of the 3rd International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS'05)*, volume 3829 of *Lecture Notes in Computer Science*, pages 2–16. Springer-Verlag, 2005.

[MNP06] Oded Maler, Dejan Nickovic, and Amir Pnueli. From MITL to timed automata. In Eugene Asarin and Patricia Bouyer, editors, *Proceedings of the 4th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS'06)*, volume 4202 of *Lecture Notes in Computer Science*, pages 274–289. Springer-Verlag, 2006.

[OW05] Joël Ouaknine and James Worrell. On the decidability of metric temporal logic. In *Proceedings of the 20th Annual IEEE Symposium on Logic in Computer Science (LICS'05)*, pages 188–197. IEEE Computer Society Press, 2005.

[OW07] Joël Ouaknine and James Worrell. On the decidability and complexity of metric temporal logic over finite words. *Logical Methods in Computer Science*, 3(1), 2007.

[Rab07] Alexander Rabinovich. Temporal logics with incommensurable distances are undecidable. *Information and Computation*, 205(5):707–715, 2007.

[Rab08] Alexander Rabinovich. Complexity of metric temporal logic with counting. In Franck Cassez and Claude Jard, editors, *Proceedings of the 6th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS'08)*, Lecture Notes in Computer Science. Springer-Verlag, 2008.

[Wil94] Thomas Wilke. Specifying timed state sequences in powerful decidable logics and timed automata. In Hans Langmaack, Willem P. de Roever, and Jan Vytopil, editors, *Proceedings of the 3rd International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT'94)*, volume 863 of *Lecture Notes in Computer Science*, pages 694–715. Springer-Verlag, 1994.