MTL with Bounded Variability: Decidability and Complexity

Carlo A. Furia and Matteo Rossi

Dipartimento di Elettronica e Informazione, Politecnico di Milano, Italy
{furia, rossi}@elet.polimi.it
http://home.dei.polimi.it/lastname/

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.

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 [2].[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 [2]. The other adopts the syntactic concession of restricting MTL formulas to a subset known as MITL [1]. More recently other syntactic adjustments have been studied [4].

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

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

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 [11]. 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 [5] 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 [17]). For lack of space, several proofs and details are omitted from the paper but can be found in [9].

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 [2]. In [2] 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 [14]. 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 [1], 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 [1] is based on a complex translation into timed automata; similar, but simpler, automata-based techniques have been studied by Maler et al. [13].

Hirshfeld and Rabinovich have reconsidered the work on MITL from a broader, more foundational, perspective built upon the standard timed behavior model [11]. 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 [10], 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) [16].

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. [4] 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. [12].

Dense-timed words where the maximum number of events in a unit interval is fixed have been introduced by Wilke in [17]. 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 [6] (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.

In previous work [8, 7], we introduced the notion of non-Berkeleyness:[2] 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 $t \in \mathbb{S}$. Correspondingly, we define the *length* $|\mathbf{x}|$ of $\mathbf{x}$ as $|\mathbf{x}| = n$.

2.1 Words and Behaviors

The two most popular models of real-time behavior [2, 3] are the *timed word* (also called timed state sequence [17]) and the *timed behavior* (also called Boolean signal [13], timed interval sequence [1], or trajectory [6]). 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.

[2] The name is an analogy with non-Zenoness [7].

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* [11].

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) $\mathbb{I}$ 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) the intervals in $\iota(b)$ form a partition of $\mathbb{T}$; (3) for all $i \in \mathbb{I}$ we have $\tau_i = u_i = l_{i+1}$; (4) 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.

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.

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*.

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 $\mathcal{B}\Sigma\mathbb{T}_{k,\delta}$. Formally, $\mathcal{B}\Sigma\mathbb{T}_{k,\delta} = \{b \in \mathcal{B}\Sigma\mathbb{T} \mid \forall t \in \mathbb{T} : |[t, t+\delta] \cap \tau(b)| \leq k\}$.

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. $\mathcal{T}\Sigma\mathbb{T}^\omega_{\exists k \exists \delta} = \bigcup_{\substack{k \in \mathbb{N}_{>0} \\ \delta \in \mathbb{R}_{>0}}} \mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta}$, $\mathcal{T}\Sigma\mathbb{T}^*_{\exists k \exists \delta} = \bigcup_{\substack{k \in \mathbb{N}_{>0} \\ \delta \in \mathbb{R}_{>0}}} \mathcal{T}\Sigma\mathbb{T}^*_{k,\delta}$).

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. $\mathcal{T}\Sigma\mathbb{T}^\omega_{\exists\delta} = \bigcup_{\delta \in \mathbb{R}_{>0}} \mathcal{T}\Sigma\mathbb{T}^\omega_\delta$, $\mathcal{T}\Sigma\mathbb{T}^*_{\exists\delta} = \bigcup_{\delta \in \mathbb{R}_{>0}} \mathcal{T}\Sigma\mathbb{T}^*_\delta$).

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* $k > k' \geq 2$*:*

$$\mathcal{B}\Sigma\mathbb{T}_{1,\delta'} \subset \mathcal{B}\Sigma\mathbb{T}_{\delta'} \subset \mathcal{B}\Sigma\mathbb{T}_\delta \subset \mathcal{B}\Sigma\mathbb{T}_{k',\delta} \subset \mathcal{B}\Sigma\mathbb{T}_{k,\delta} \subset \mathcal{B}\Sigma\mathbb{T}_{\exists k \exists \delta} \subset \mathcal{B}\Sigma\mathbb{T} \tag{1}$$

$$\mathcal{B}\Sigma\mathbb{T}_\delta \subset \mathcal{B}\Sigma\mathbb{T}_{\exists\delta} \subset \mathcal{B}\Sigma\mathbb{T}_{\exists k \exists \delta} \subset \mathcal{B}\Sigma\mathbb{T} \tag{2}$$

$$\mathcal{B}\Sigma\mathbb{T}_{\exists\delta} \textit{ and } \mathcal{B}\Sigma\mathbb{T}_{k',\delta} \textit{ are incomparable} \tag{3}$$

$$\mathcal{T}\Sigma\mathbb{T}^\omega_{\delta'} \subset \mathcal{T}\Sigma\mathbb{T}^\omega_\delta = \mathcal{T}\Sigma\mathbb{T}^\omega_{1,\delta} \subset \mathcal{T}\Sigma\mathbb{T}^\omega_{k',\delta} \subset \mathcal{T}\Sigma\mathbb{T}^\omega_{k,\delta} \subset \mathcal{B}\Sigma\mathbb{T}_{\exists k \exists \delta} \subset \mathcal{T}\Sigma\mathbb{T}^\omega \tag{4}$$

$$\mathcal{T}\Sigma\mathbb{T}^\omega_\delta \subset \mathcal{T}\Sigma\mathbb{T}^\omega_{\exists\delta} \subset \mathcal{B}\Sigma\mathbb{T}_{\exists k \exists \delta} \subset \mathcal{T}\Sigma\mathbb{T}^\omega \tag{5}$$

$$\mathcal{T}\Sigma\mathbb{T}^\omega_{\exists\delta} \textit{ and } \mathcal{T}\Sigma\mathbb{T}^\omega_{k',\delta} \textit{ are incomparable} \tag{6}$$

$$\mathcal{T}\Sigma\mathbb{T}^*_{\delta'} \subset \mathcal{T}\Sigma\mathbb{T}^*_\delta = \mathcal{T}\Sigma\mathbb{T}^*_{1,\delta} \subset \mathcal{T}\Sigma\mathbb{T}^*_{k',\delta} \subset \mathcal{T}\Sigma\mathbb{T}^*_{k,\delta} \subset \mathcal{T}\Sigma\mathbb{T}^*_{\exists\delta} = \mathcal{T}\Sigma\mathbb{T}^*_{\exists k \exists \delta} = \mathcal{T}\Sigma\mathbb{T}^* \tag{7}$$

For lack of space, in the rest of the paper we will focus on behaviors, leaving the proofs of the corresponding results for words to [9]. Also, 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 [15]. 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 [1].

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

Let us start with the Metric Interval Temporal Logic (MITL) [1], 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: $\psi ::= \mathsf{p} \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \mathsf{U}_I(\psi_1, \psi_2) \mid \mathsf{S}_I(\psi_1, \psi_2)$. Metric Temporal Logic (MTL) [2] is defined simply as an extension of MITL where singular intervals are allowed.

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: $\Diamond_I(\phi) = \mathsf{U}_I(\top, \phi)$, $\Box_I(\phi) = \neg\Diamond_I(\neg\phi)$, $\mathsf{R}_I(\phi_1, \phi_2) = \neg\mathsf{U}_I(\neg\phi_1, \neg\phi_2)$, $\bigcirc(\phi) = \mathsf{U}(\phi, \top)$, as well as their past counterparts $\overleftarrow{\Diamond}_I(\phi) = \mathsf{S}_I(\top, \phi)$, $\overleftarrow{\Box}_I(\phi) = \neg\overleftarrow{\Diamond}_I(\neg\phi)$, $\mathsf{T}_I(\phi_1, \phi_2) = \neg\mathsf{S}_I(\neg\phi_1, \neg\phi_2)$, $\overleftarrow{\bigcirc}(\phi) = \mathsf{S}(\phi, \top)$, and $\mathrm{Alw}(\phi) = \overleftarrow{\Box}(\phi) \wedge \phi \wedge \Box(\phi)$ and $\triangle(\phi) = \overleftarrow{\bigcirc}(\neg\phi) \wedge (\phi \vee \bigcirc(\phi))$.

Semantics. For $b \in \overline{\mathcal{B}\Sigma\mathbb{T}}$ (with $\Sigma = 2^{\mathcal{P}}$) and $t \in \mathbb{T}$ we define:[3]

$b(t) \models_{\mathbb{T}} \mathsf{p}$	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	$b(0) \models_{\mathbb{T}} \phi$

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

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: $\mathsf{p} \mid \neg\psi \mid \psi_1 \wedge \psi_2 \mid \mathsf{U}(\psi_1, \psi_2) \mid \mathsf{S}(\psi_1, \psi_2) \mid \Diamond_{<\mathrm{d}}(\psi) \mid \overleftarrow{\Diamond}_{<\mathrm{d}}(\psi)$. The fact that every MITL formula can be expressed according to the syntax above follows from two results. [11, 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, [1, Lm. 2.16] showed that every MITL using intervals with rational endpoints can be translated into an equisatisfiable 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: $\Diamond_{=\mathrm{d}}(\phi) \mid \overleftarrow{\Diamond}_{=\mathrm{d}}(\phi)$.

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:

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

Also, 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 QITL: Decidable Extensions of MITL

Following [11, 16], we introduce decidable extensions of MITL over behaviors. These extensions will be useful in the decidability proofs of Section 5.

We extend MITL by introducing modalities $\diamondsuit_I^n(\psi_1, \ldots, \psi_n)$ for $n > 0$ and I a non-singular interval. We denote the corresponding temporal logics by $\mathrm{QITL}(n)$, for $n > 0$. Also, we denote the temporal logic $\bigcup_{k>0} \mathrm{QITL}(k)$ simply by QITL; note that QITL is essentially equivalent to the logic TLPI introduced in [16]. The semantics of the new operators is:

$b(t) \models_{\mathbb{T}} \diamondsuit_I^n(\psi_1, \ldots, \psi_n)$ 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 $\diamondsuit_I^n(\psi) = \diamondsuit_I^n\left(\underbrace{\psi, \psi, \ldots, \psi}_{n \text{ times}}\right)$.

The small syntactic gap between QITL and TLPI can be easily bridged along the lines of [11] (see [9] for details). Hence, the following is a corollary of the complexity results for TLPI — with a succinct encoding of constants — presented in [16].

Proposition 2 (Decidability and Complexity of QITL**).** QITL *is decidable with an* ***EXPSPACE****-complete validity problem.*

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.

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

Note that $\theta(\phi) = \mathrm{O}(|\phi|)$. Then, we can prove the following.

Lemma 1. *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; for brevity let us just consider the cases $\phi = \Diamond_{=\mathrm{d}}(\phi')$ and $\phi = \Diamond_{<\mathrm{d}}(\phi')$.

In the first case, clearly $\tau(\phi) = \ldots, x_{-1} - \mathrm{d}, x_0 - \mathrm{d}, x_1 - \mathrm{d}, \ldots, x_i - \mathrm{d}, \ldots$, where $\tau(\phi') = \ldots, x_{-1}, x_0, x_1, \ldots, x_i, \ldots$. Thus, $b_{\phi'} \in \mathcal{B}\Sigma\mathbb{T}_{\theta(\phi'),\delta}$ implies $b_\phi \in \mathcal{B}\Sigma\mathbb{T}_{\theta(\phi),\delta}$ as well, since $\theta(\phi) = \theta(\phi')$.

For the second case, 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 $b_\phi \in \mathcal{B}\Sigma\mathbb{T}_{\theta(\phi),\delta}$. 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

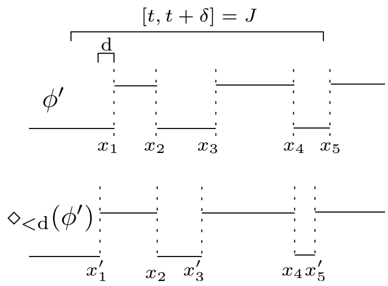

Fig. 1. $b_{\phi'}$ and $b_{\Diamond_{<\mathrm{d}}(\phi')}$ over J.

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. □

4.2 Describing Sequences of Transitions

Let us now introduce QITL formula $\mathsf{happ}(\phi, k, I)$ stating that formula ϕ takes exactly $k-1$ consecutive transitions, eventually leading to true (i.e., it holds at the end of I). Formally, 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.

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, nonsingular interval I, and $n > 0$, we introduce the QITL formula:

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

Lemma 2. *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 one can check that $b(t) \models \mathsf{happ}(\phi, k+1, [\tau - \delta, \tau])$, 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 $b(t+\tau) \models \phi$. 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 one can check that $b(t) \models \mathsf{happ}(\phi, h+1, [0, \tau])$, 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 $b(t+\tau) \models \phi$. $\square$

4.3 Syntactic Characterizations

This section defines non-Berkeleyness and bounded variability syntactically.

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

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

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 $|\mathcal{P}|$. To this end, let us first define: $\mathsf{RT}(\beta) = \triangle(\beta) \wedge \beta \vee \triangle(\neg\beta) \wedge \neg\beta$, $\mathsf{LT}(\beta) = \triangle(\beta) \wedge \neg\beta \vee \triangle(\neg\beta) \wedge \beta$, $\mathsf{GT}(\beta) = \triangle(\beta) \vee \triangle(\neg\beta)$ that model a right-continuous, left-continuous, and generic transition of β, respectively. Then, we introduce:

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

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

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

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

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 $b \in \mathcal{B}\Sigma\mathbb{T}_\delta$ with $\Sigma = 2^{\mathcal{P}}$ iff $b \models \chi'_\delta$, and χ'_δ has size quadratic in $|\mathcal{P}|$.

Bounded variability. 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 $b \in \mathcal{B}\Sigma\mathbb{T}_{k,\delta}$ 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.

Lemma 3. *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 1, ϕ undergoes at most $\theta(\phi)$ transitions over $t \oplus I$. So, from Lemma 2, we have immediately that $b(t+\mathrm{d}) \models \phi$ — i.e., $b(t) \models \Diamond_{=\mathrm{d}}(\phi)$ — iff $b(t) \models \mathsf{yieldsT}(\phi, \theta(\phi)+1, I)$. □

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.

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

Theorem 1. *For any MTL formula ϕ, for any behavior $b \in \mathcal{B}\Sigma\mathbb{T}_\delta$ for some $\delta > 0$, we have $b \models_{\mathbb{T}} \phi$ iff $b \models_{\mathbb{T}} \mu(\phi)$.*

Theorem 1, the decidability of MITL and QITL [1, 11], and the syntactic characterization of non-Berkeleyness by means of the χ_δ formula, immediately imply the following.

Corollary 1. *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 1, 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 2. □

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 1, 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 $\mathcal{B}\Sigma\mathbb{T}_{k,\delta}$.

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. A partial negative answer to this question can be provided by showing that MITL is strictly less expressive than MTL over non-Berkeley behaviors. 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. However, one can prove that for any $\delta > 0$, MTL is strictly more expressive than MITL over $\mathcal{B}\Sigma\mathbb{T}_\delta$. We refer to [9] for some details of the (involved) proof.

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 2. *The satisfiability problem for MTL over $\mathcal{B}\Sigma\mathbb{T}_\delta$ is **EXPSPACE**-complete.*

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 2).

The **EXPSPACE**-hardness of MTL satisfiability over non-Berkeley behaviors can be proved by reducing the corresponding problem over the integers, where integer-timed words are embedded into non-Berkeley behaviors. See [9] for details. □

With a very similar justification we can prove the following.

Theorem 3. *The satisfiability problem for MTL over $\mathcal{B}\Sigma\mathbb{T}_{k,\delta}$ is **EXPSPACE**-complete (assuming a unary encoding of k).*

7 Undecidability Results

MTL is decidable no more if we consider all non-Berkeley behaviors for any δ together. More precisely, the satisfiability problem for MTL over $\mathcal{B}\Sigma\mathbb{T}_{\exists\delta}$ is Σ_1^0-complete; compare against the same problem over $\mathcal{B}\Sigma\mathbb{T}$ where it is Σ_1^1-complete.

Theorem 4. *The satisfiability problem for MTL over $\mathcal{B}\Sigma\mathbb{T}_{\exists\delta}$ is $\Sigma_1^0 =$ **RE**-complete.*

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

To show **RE**-hardness, we reduce the halting problem for 2-counter machines to MTL satisfiability over $\mathcal{B}\Sigma\mathbb{T}_{\exists\delta}$. The key insight is that a halting computation is one where only a finite number of instructions is executed. Correspondingly it can be represented by a behavior where only a finite number of transitions occur within a finite amount of time; such behaviors are necessarily in $\mathcal{B}\Sigma\mathbb{T}_{\exists\delta}$ because the infimum over distances between transitions coincides with the minimum. See [9] for all details. □

The above proof can be adapted with simple modifications to work for infinite timed words, as well as for the classes $\mathcal{B}\Sigma\mathbb{T}_{\exists k\exists\delta}$ and $\mathcal{T}\Sigma\mathbb{T}^{\omega}_{\exists k\exists\delta}$. On the contrary, undecidability does not carry over to finite words, where the problem is known to be decidable [14].

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.

	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^0_1-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^0_1-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.

Acknowledgements. We thank Alexander Rabinovich for providing us with a draft of his recent work [16] and the anonymous reviewers of FORMATS for their detailed and insightful remarks.

References

1. R. Alur, T. Feder, and T. A. Henzinger. The benefits of relaxing punctuality. *Journal of the ACM*, 43(1):116–146, 1996.
2. R. Alur and T. A. Henzinger. Real-time logics: Complexity and expressiveness. *Information and Computation*, 104(1):35–77, 1993.
3. E. Asarin, P. Caspi, and O. Maler. Timed regular expressions. *Journal of the ACM*, 49(2):172–206, 2002.
4. P. Bouyer, N. Markey, J. Ouaknine, and J. Worrell. The cost of punctuality. In *Proceedings of LICS'07*, 2007.
5. D. D'Souza and P. Prabhakar. On the expressiveness of MTL in the pointwise and continuous semantics. *Formal Methods Letters*, 9(1):1–4, 2007.
6. M. Fränzle. Decidability of duration calculi on restricted model classes. Technical Report Kiel MF 21/1, Christian-Albrechts Universität Kiel, 1996.
7. C. A. Furia, M. Pradella, and M. Rossi. Automated verification of dense-time MTL specifications via discrete-time approximation. In *Proceedings of FM'08*, volume 5014 of *LNCS*, pages 132–147, 2008.
8. C. A. Furia and M. Rossi. Integrating discrete- and continuous-time metric temporal logics through sampling. In *Proceedings of FORMATS'06*, volume 4202 of *LNCS*, pages 215–229, 2006.
9. C. A. Furia and M. Rossi. MTL with bounded variability: Decidability and complexity. Technical report, DEI, Politecnico di Milano, May 2008.
10. D. M. Gabbay, I. Hodkinson, and M. Reynolds. *Temporal Logic (vol. 1)*. Oxford University Press, 1994.
11. Y. Hirshfeld and A. M. Rabinovich. Logics for real time: Decidability and complexity. *Fundamenta Informaticae*, 62(1):1–28, 2004.
12. C. Lutz, D. Walther, and F. Wolter. Quantitative temporal logics over the reals: PSPACE and below. *Information and Computation*, 205(1):99–123, 2007.
13. O. Maler, D. Nickovic, and A. Pnueli. From MITL to timed automata. In *Proceedings of FORMATS'06*, volume 4202 of *LNCS*, pages 274–289, 2006.
14. J. Ouaknine and J. Worrell. On the decidability and complexity of metric temporal logic over finite words. *Logical Methods in Computer Science*, 3(1), 2007.
15. A. Rabinovich. Temporal logics with incommensurable distances are undecidable. *Information and Computation*, 205(5):707–715, 2007.
16. A. M. Rabinovich. Complexity of metric temporal logic with counting. In *Proceedings of FORMATS'08*, LNCS, 2008.
17. T. Wilke. Specifying timed state sequences in powerful decidable logics and timed automata. In *Proceedings of FTRTFT'94*, volume 863 of *LNCS*, pages 694–715, 1994.