

A Theory of Sampling for Continuous-Time Metric Temporal Logic

Carlo A. Furia

ETH Zurich

and

Matteo Rossi

Politecnico di Milano

This paper revisits the classical notion of sampling in the setting of real-time temporal logics for the modeling and analysis of systems. The relationship between the satisfiability of Metric Temporal Logic (MTL) formulas over continuous-time models and over discrete-time models is studied. It is shown to what extent discrete-time sequences obtained by sampling continuous-time signals capture the semantics of MTL formulas over the two time domains. The main results apply to “flat” formulas that do not nest temporal operators and can be applied to the problem of reducing the verification problem for MTL over continuous-time models to the same problem over discrete-time, resulting in an automated partial practically-efficient discretization technique.

Categories and Subject Descriptors: D.2.4 [Software Engineering]: Software/Program Verification; F.3.1 [Logics and Meanings of Programs]: Specifying and Verifying and Reasoning about Programs; F.4.3 [Mathematical Logic and Formal Languages]: Formal Languages

General Terms: Theory, Verification

Additional Key Words and Phrases: Real-time, metric temporal logic, hybrid systems

1. INTRODUCTION

Computer programs are inherently *discrete* items, and they are typically modeled through techniques from the discrete mathematics domain. If, however, one shifts from a computer-centric to a *system-centric* view [Furia et al. 2010], physical elements, which are best described through *continuous* signals, enter the picture and must be taken into account throughout the system development process. This is the challenge that is at the core of the research on real-time and hybrid systems [Henzinger and Sifakis 2006]. The challenge has two facets: *modeling* systems that integrate continuous and discrete components and *analyzing* properties of the integrated systems.

Author’s address: carlo.furia@inf.ethz.ch.

A preliminary version of this work appeared in [Furia and Rossi 2006; Furia et al. 2008a].

This work has been partially supported by the Italian Government under the project ArtDeco (FIRB RBNE05C3AH) and by the European Commission, Programme IDEAS-ERC, Project 227977-SMScom.

Permission to make digital/hard copy of all or part of this material without fee for personal or classroom use provided that the copies are not made or distributed for profit or commercial advantage, the ACM copyright/server notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists requires prior specific permission and/or a fee.

© 20YY ACM 1529-3785/20YY/0700-0001 \$5.00

In this article we develop some techniques for the modeling and analysis of real-time systems with mixed continuous- and discrete-time components. Our approach targets the well-known Metric Temporal Logic (MTL [Koymans 1990; Alur and Henzinger 1993]) as formal notation, and it is based on the classical notion of *sampling*.

Sampling is a widely-used technique in the engineering domain, in particular in signal processing and automatic control, whereby continuous-time signals are transformed in discrete-time counterparts that are more amenable to digital processing [Benedetto and Ferreira 2001]. In systems where continuous- and discrete-time components interact, a *sampler* constitutes the interface between these two classes of components, as it retains some *partial* information of the continuous-time processes and passes it to the discrete-time parts (see Figure 1). The classical sampling theory




Fig. 1. A system with a sampler.

determines qualitatively how much information is preserved in this discretization process, and when the continuous-time signal can be perfectly reconstructed solely from its discrete-time samplings.

The sampling approach described in this article borrows from these well-known ideas, but revisits them in the very different setting of formal modeling and analysis of systems with real-time temporal logics.¹

In our approach, the behavior of components is modeled by means of MTL formulas. MTL formulas can be given a continuous-time or discrete-time semantics by interpreting them over sets of continuous- or discrete-time *behaviors*² respectively (see Section 2 for formal definitions). Accordingly, MTL formulas can model either continuous- or discrete-time components. The problem of providing a unified semantics is then solved by introducing simple *syntactic transformations* to be applied to MTL formulas when moving their interpretation from continuous to discrete time, or *vice versa*. These transformations take into account the information that is preserved under sampling. That is, given a continuous-time formula $\phi_{\mathbb{R}}$, its (transformed) discrete-time counterpart $\phi_{\mathbb{Z}}$ is satisfied precisely by all discrete-time behaviors that are obtained by sampling the continuous-time behaviors satisfying $\phi_{\mathbb{R}}$. Information preservation requires however an additional requirement — called non-Berkeleyness — on the sampled continuous-time behaviors to ensure that they are sufficiently “slow” with respect to the speed of the sampling process.

In summary, the contribution of this article is twofold. First, it introduces conditions that allow us to precisely relate the satisfiability of continuous-time MTL formulas to that of some suitable, “sampled”, discrete-time counterparts. Second, it

¹Section 5.3 discusses in some detail how the classical notion of sampling and the one presented here are related, though different.

²Also called (Boolean) *signals* [Maler et al. 2006; Hirshfeld and Rabinovich 2004].

exploits this relation to define an effective, albeit partial, automated analysis technique that can be used to prove (or disprove) properties of systems with continuous-time components by reduction to the (usually simpler) discrete-time case. In this paper we do not deal with aspects regarding its implementation and performance in practice, which have been dealt with in related work [Furia et al. 2008a; 2008b; Bersani et al. 2009]. Rather, we focus on the mathematical concepts underlying the relation between continuous- and discrete-time semantics of MTL.

This article is structured as follows. Section 2 introduces the MTL notation and its formal semantics, and discusses the expressiveness of some of its significant subsets. Section 3 presents the notions of sampling and sampling invariance for MTL, and proves some fundamental results about significant subsets of the MTL language that are amenable to the sampling technique introduced beforehand, and hence are suitable to define a unified semantics. Section 4 shows how the results of Section 3 can be applied to the problem of automated verification of continuous-time systems described with MTL. Finally, Section 5 provides an overview of related work, focusing on a few well-known approaches that are similar to ours; Section 6 briefly concludes.

Let us remark that the mathematical distinction between continuous and merely dense time models does not impact the results of this paper. Accordingly, we will essentially use the two terms as synonyms.

2. METRIC TEMPORAL LOGIC(S)

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 any numbers x, y with $y > 0$, $x \pm \infty/y$ is defined to be $\pm\infty$. We occasionally represent intervals by pseudo-arithmetic expressions such as $> x$, $\geq x$, $< x$, $\leq x$, and $= x$ for (x, ∞) , $[x, \infty)$, $[0, x)$, $[0, x]$ and $[x, x]$, respectively. For simplicity, we sometimes relax the notation for unbounded intervals and represent them with square — rather than round — closing brackets.

2.1 Behaviors

In this paper, \mathbb{T} denotes any of the two time domains \mathbb{R} and \mathbb{Z} . It is not difficult to adapt most notions and results to their mono-infinite counterparts $\mathbb{R}_{\geq 0}$ and \mathbb{N} , and possibly to other dense and discrete sets suitable to represent time domains [Koymans 1992]. Also, let \mathcal{P} be a set of propositional letters.

DEFINITION 1 BEHAVIORS. *A (timed) behavior over time domain \mathbb{T} and alphabet \mathcal{P} is a function $b : \mathbb{T} \rightarrow 2^{\mathcal{P}}$ which maps every time instant $t \in \mathbb{T}$ to the set of propositions $b(t) \in 2^{\mathcal{P}}$ that hold at t . The set of all behaviors over time domain \mathbb{T} and alphabet \mathcal{P} is denoted by $\overline{\mathcal{BPT}}$.*

$b|_P$ is a behavior over alphabet $P \subseteq \mathcal{P}$, denoting the projection of b over P . For a behavior b over some *dense* time domain \mathbb{T} , let $\tau(b)$ denote the ordered (multi)set of its discontinuity points, that is $\tau(b) = \{x \in \mathbb{T} \mid b(x) \neq \lim_{t \rightarrow x^-} b(t), \text{ or } b(x) \neq \lim_{t \rightarrow x^+} b(t), \text{ or any of the two limits does not exist}\}$, where each point that is both a right- and a left-discontinuity appears twice in $\tau(b)$. When \mathbb{T} is a discrete set, $\tau(b)$ is defined to be the time domain \mathbb{T} itself. If $\tau(b)$ is discrete, we can represent it as an ordered sequence (possibly unbounded to $\pm\infty$) of elements τ_i for $i \in \mathbb{I}$; 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.

2.1.0.1 *Non-Zenoness*. Since one is typically interested only in behaviors that represent physically meaningful behaviors, it is common to assume some regularity requirements. In particular, it is customary to assume *non-Zenoness*, also called *finite variability* [Hirshfeld and Rabinovich 2004].

DEFINITION 2 NON-ZENONESS. A behavior $b \in \overline{\mathcal{BPT}}$ is non-Zeno iff $\tau(b)$ has no accumulation points. The set of all non-Zeno behaviors is denoted by \mathcal{BPT} .

Notice that discrete-time behaviors are trivially non-Zeno. Also, 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 $b \in \mathcal{BPT}$, $\iota(b)$ is an ordered sequence of intervals $\iota(b) = \{I_i = \langle^i l_i, u_i \rangle^i\}$ for $i \in \mathbb{I}$ 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 set $\tau(b)$ 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.

2.1.0.2 *Non-Berkeleyness*. Some of the results of this paper will require a stronger regularity requirement than non-Zenoness, named “non-Berkeleyness” [Furia et al. 2008a].

DEFINITION 3 NON-BERKELEYNESS. A behavior $b \in \mathcal{BPT}$ 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{BPT} that are non-Berkeley for δ is denoted by \mathcal{BPT}_δ ; with the notation introduced above, it is $\mathcal{BPT}_\delta = \{b \in \mathcal{BPT} \mid \forall I \in \iota(b) : \exists t \in I : [t, t + \delta] \subseteq I\}$. A behavior that is not non-Berkeley for every positive δ is called Berkeley.

Any behavior where some proposition holds at an isolated point t is Berkeley: any $\delta > 0$ is such that $[t, t + \delta] \not\subseteq [t, t]$.

2.2 MTL: Syntax and Semantics

This section defines formally the syntax and semantics of MTL.

2.2.1 *MTL Syntax.* In this paper only *propositional* temporal logics are considered; correspondingly, the elementary building block of temporal logic formulas is defined.

DEFINITION 4 PROPOSITIONAL FORMULAS. *Propositional formulas* $\pi \in \text{PL}$ are defined by the grammar $\pi ::= \mathbf{p} \mid \neg \mathbf{p} \mid \pi_1 \wedge \pi_2 \mid \pi_1 \vee \pi_2$ — for $\mathbf{p} \in \mathcal{P}$ — as Boolean combinations of propositional letters.

MTL formulas are obtained by combining propositional formulas with the *bounded until* \mathbf{U}_I metric modality, as well as its past counterpart *bounded since* \mathbf{S}_I . We assume a *negation normal form* (NNF) syntax, where negations are pushed down to atomic propositions, as this will simplify the presentation of the results. Correspondingly, *bounded release* \mathbf{R}_I and *bounded trigger* \mathbf{T}_I operators — duals to the *until* and *since* operators, respectively — are introduced as primitive modalities.

DEFINITION 5 MTL FORMULAS. *MTL formulas for a time domain* \mathbb{T} are defined by the grammar:

$$\phi ::= \pi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \mathbf{U}_I(\phi_1, \phi_2) \mid \mathbf{S}_I(\phi_1, \phi_2) \mid \mathbf{R}_I(\phi_1, \phi_2) \mid \mathbf{T}_I(\phi_1, \phi_2)$$

where $\pi \in \text{PL}$ ranges over propositional formulas and I ranges over (possibly unbounded) intervals of the time domain \mathbb{T} with endpoints in $\mathbb{T} \cap \mathbb{Q} \cup \{\pm\infty\}$ (notice that negative endpoints are allowed).

Henceforth, we will drop interval I in modalities when it is $[0, +\infty)$.

The results of this paper are focused on the *flat* subset $\mathfrak{b}\text{MTL}$ of MTL, whose formulas do not nest temporal operators.³

DEFINITION 6 FLAT MTL FORMULAS. $\mathfrak{b}\text{MTL}$ formulas for a time domain \mathbb{T} are defined by the grammar:

$$\phi ::= \pi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \vee \phi_2 \mid \mathbf{U}_I(\pi_1, \pi_2) \mid \mathbf{S}_I(\pi_1, \pi_2) \mid \mathbf{R}_I(\pi_1, \pi_2) \mid \mathbf{T}_I(\pi_1, \pi_2)$$

where $\pi, \pi_1, \pi_2 \in \text{PL}$ range over propositional formulas and I ranges over (possibly unbounded) intervals of the time domain \mathbb{T} with endpoints in $\mathbb{T} \cap \mathbb{Q} \cup \{\pm\infty\}$.

In the remainder of the paper, the following other MTL subsets will be needed.

- LTL is the MTL subset where all intervals I are $[0, +\infty)$ (i.e., all operators are *qualitative*).
- Υ - $\mathfrak{b}\text{MTL}$, with Υ any given set of MTL formulas, is the MTL subset defined by the same grammar as $\mathfrak{b}\text{MTL}$, except that π is allowed to range over $\text{PL} \cup \Upsilon$.
- An MTL formula is *discrete-endpoint* if all its intervals have endpoints in $\mathbb{Z} \cup \{\pm\infty\}$.
- An MTL formula is *dense-endpoint* if all its intervals have endpoints in $\mathbb{R} \cup \{\pm\infty\}$. It is clear that any MTL formula is dense-endpoint; we will use this redundant terminology whenever useful to characterize formulas to be interpreted over a dense time domain, as opposed to a discrete one.

³Different notions of flatness for (metric) temporal logic have been introduced in the literature [Dams 1999; Comon and Cortier 2000; Bouyer et al. 2007].

2.2.2 *MTL Semantics.* We define MTL semantics parametrically with respect to the time domain \mathbb{T} .

DEFINITION 7 MTL SEMANTICS. *Let $b \in \mathcal{BP}\mathbb{T}$ be a behavior over \mathcal{P} and time domain \mathbb{T} . For $t \in \mathbb{T}$, MTL semantics is defined recursively as follows.⁴*

$b(t) \models_{\mathbb{T}} \mathbf{p}$	iff	$\mathbf{p} \in b(t)$
$b(t) \models_{\mathbb{T}} \neg \mathbf{p}$	iff	$\mathbf{p} \notin b(t)$
$b(t) \models_{\mathbb{T}} \phi_1 \wedge \phi_2$	iff	$b(t) \models_{\mathbb{T}} \phi_1$ and $b(t) \models_{\mathbb{T}} \phi_2$
$b(t) \models_{\mathbb{T}} \phi_1 \vee \phi_2$	iff	$b(t) \models_{\mathbb{T}} \phi_1$ or $b(t) \models_{\mathbb{T}} \phi_2$
$b(t) \models_{\mathbb{T}} \mathbf{U}_I(\phi_1, \phi_2)$	iff	$\exists d \in I$ s.t.: $t + d \in \mathbb{T}$, $b(t + d) \models_{\mathbb{T}} \phi_2$, and $\forall t' \in [0, d) \oplus t \cap \mathbb{T}$ it is $b(t') \models_{\mathbb{T}} \phi_1$
$b(t) \models_{\mathbb{T}} \mathbf{S}_I(\phi_1, \phi_2)$	iff	$\exists d \in I$ s.t.: $t - d \in \mathbb{T}$, $b(t - d) \models_{\mathbb{T}} \phi_2$, and $\forall t' \in -[0, d) \oplus t \cap \mathbb{T}$ it is $b(t') \models_{\mathbb{T}} \phi_1$
$b(t) \models_{\mathbb{T}} \mathbf{R}_I(\phi_1, \phi_2)$	iff	$\forall d \in I$ s.t. $t + d \in \mathbb{T}$ it is: $b(t + d) \models_{\mathbb{T}} \phi_2$ or $\exists t' \in [0, d) \oplus t \cap \mathbb{T}$ s.t. $b(t') \models_{\mathbb{T}} \phi_1$
$b(t) \models_{\mathbb{T}} \mathbf{T}_I(\phi_1, \phi_2)$	iff	$\forall d \in I$ s.t. $t - d \in \mathbb{T}$ it is: $b(t - d) \models_{\mathbb{T}} \phi_2$ or $\exists t' \in -[0, d) \oplus t \cap \mathbb{T}$ s.t. $b(t') \models_{\mathbb{T}} \phi_1$

If $b(t) \models_{\mathbb{T}} \phi$ holds for all $t \in \mathbb{T}$ we write $b \models_{\mathbb{T}} \phi$.

We denote by $\llbracket \phi \rrbracket_{\mathbb{T}}$ (respectively $\llbracket \phi \rrbracket_{\mathbb{T}}^{\delta}$) the set of all non-Zeno (respectively non-Berkeley for δ) models of formula ϕ over \mathbb{T} , i.e., $\llbracket \phi \rrbracket_{\mathbb{T}} \triangleq \{b \in \mathcal{BP}\mathbb{T} \mid b \models_{\mathbb{T}} \phi\}$ (respectively $\llbracket \phi \rrbracket_{\mathbb{T}}^{\delta} \triangleq \{b \in \mathcal{BP}\mathbb{T}_{\delta} \mid b \models_{\mathbb{T}} \phi\}$). If $\llbracket \phi \rrbracket_{\mathbb{T}}$ is empty, ϕ is called \mathbb{T} -unsatisfiable, and \mathbb{T} -satisfiable otherwise. If $\llbracket \phi \rrbracket_{\mathbb{T}}$ coincides with $\mathcal{BP}\mathbb{T}$, ϕ is called \mathbb{T} -valid. Similar definitions are assumed for \mathbb{T}^{δ} -satisfiability and \mathbb{T}^{δ} -validity, with respect to $\llbracket \phi \rrbracket_{\mathbb{T}}^{\delta}$. For $b \in \mathcal{BP}\mathbb{T}$, we define the derived behavior b_{ϕ} that represents the truth value of ϕ over b as:

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

For propositional letters \mathbf{a}, \mathbf{b} , $b^{\mathbf{a} \setminus \mathbf{b}}$ denotes the behavior obtained from b by renaming \mathbf{a} into \mathbf{b} .

Notice that MTL is closed under complement, even if this is not apparent in the definition of its syntax. More precisely, one can check that $b(t) \not\models_{\mathbb{T}} \mathbf{U}_I(\phi_1, \phi_2)$ holds if and only if $b(t) \models_{\mathbb{T}} \mathbf{R}_I(\neg \phi_1, \neg \phi_2)$ does, thus providing an indirect definition of negation. A similar relation holds for *since* with respect to *trigger*.

Definition 7 considers the basic modalities in their *non-strict* versions as, for instance, $\mathbf{U}(\phi_1, \phi_2)$ requires ϕ_1 to hold at the current instant; i.e., it constrains the present as well as the strict future. Also, a *global satisfiability semantics* is assumed, where $b \models_{\mathbb{T}} \phi$ entails that ϕ holds at all time instants $t \in \mathbb{T}$. This is different than the more common *initial satisfiability semantics* $\models_{\mathbb{T}}^{\text{init}}$ where $b \models_{\mathbb{T}}^{\text{init}} \phi$ is defined as simply $b(0) \models_{\mathbb{T}} \phi$. Section 2.3 discusses the impact of these choices on expressiveness.

2.2.3 *Derived Operators and Variants.* Standard abbreviations are assumed, such as for \top , \perp , \Rightarrow , and \Leftrightarrow .

⁴In this paper, the notation $b(t) \models_{\mathbb{T}} \phi$ replaces the more common $b, t \models_{\mathbb{T}} \phi$.

It is also customary to introduce a number of derived temporal operators; those used in this paper are listed in Table I. Let us remark that the definitions of Table I do not nest temporal operators, hence they define \mathfrak{b} MTL formulas if their arguments are propositional formulas.

The first set of derived operators are the quantitative versions of the well-known *eventually* \diamond and *globally* \square modalities of classic (qualitative) linear temporal logic. On the other hand, $\text{Alw}(\phi)$ declares ϕ to hold *always*, i.e., at all time instants in the future and in the past, whereas $\text{Som}(\phi)$ declares ϕ to hold *sometimes*.

The second set of derived operators are the *nowon* \bigcirc modality and its variant Δ , with their past counterparts *uptonow* $\overleftarrow{\bigcirc}$ and $\overleftarrow{\Delta}$. Over dense-time non-Zeno behaviors, $\bigcirc(\phi)$ holds at t whenever there is a non-empty open interval $E = (0, \epsilon)$ (with $\epsilon > 0$) such that ϕ holds continuously over $t \oplus E$. On the other hand, $\Delta(\phi)$ holds at t whenever ϕ holds *nowon* or ϕ holds precisely at t . These operators are useful only over dense time, as they can be seen to be trivially equivalent to their arguments over discrete time.

Finally, the last set of derived operators introduce so-called *matching variants* [Furia and Rossi 2007] of the basic *until* and *release* modalities. For instance *matching until* $\text{U}^\downarrow(\phi_1, \phi_2)$ requires both arguments ϕ_1 and ϕ_2 to hold together at some future instant, whereas $\text{U}(\phi_1, \phi_2)$ demands only ϕ_2 to hold at some future instant. The next section discusses the impact of these variants on expressiveness.

OPERATOR	\triangleq	DEFINITION
$\diamond_I(\phi)$	\triangleq	$\text{U}_I(\top, \phi)$
$\overleftarrow{\diamond}_I(\phi)$	\triangleq	$\text{S}_I(\top, \phi)$
$\square_I(\phi)$	\triangleq	$\text{R}_I(\perp, \phi)$
$\overleftarrow{\square}_I(\phi)$	\triangleq	$\text{T}_I(\perp, \phi)$
$\text{Alw}(\phi)$	\triangleq	$\overleftarrow{\square}(\phi) \wedge \square(\phi)$
$\text{Som}(\phi)$	\triangleq	$\overleftarrow{\diamond}(\phi) \vee \diamond(\phi)$
$\bigcirc(\phi)$	\triangleq	$\text{U}_{>0}(\phi, \top) \vee (\neg\phi \wedge \text{R}_{>0}(\phi, \perp))$
$\overleftarrow{\bigcirc}(\phi)$	\triangleq	$\text{S}_{>0}(\phi, \top) \vee (\neg\phi \wedge \text{T}_{>0}(\phi, \perp))$
$\Delta(\phi)$	\triangleq	$\phi \vee \bigcirc(\phi)$
$\overleftarrow{\Delta}(\phi)$	\triangleq	$\phi \vee \overleftarrow{\bigcirc}(\phi)$
$\text{U}_I^\downarrow(\phi_1, \phi_2)$	\triangleq	$\text{U}_I(\phi_1, \phi_2 \wedge \phi_1)$
$\text{S}_I^\downarrow(\phi_1, \phi_2)$	\triangleq	$\text{S}_I(\phi_1, \phi_2 \wedge \phi_1)$
$\text{R}_I^\downarrow(\phi_1, \phi_2)$	\triangleq	$\text{R}_I(\phi_1, \phi_2 \vee \phi_1)$
$\text{T}_I^\downarrow(\phi_1, \phi_2)$	\triangleq	$\text{T}_I(\phi_1, \phi_2 \vee \phi_1)$

Table I. MTL derived temporal operators

The value of propositional formulas change at most every δ time units over non-Berkeley behaviors \mathcal{BPR}_δ ; more precisely the following holds.

LEMMA 8. *Let $b \in \mathcal{BPR}_\delta$, $t \in \mathbb{R}$, and $\pi \in \text{PL}$, such that $b(t) \models_{\mathbb{R}} \pi$. There exist $c_n, c_p \in \mathbb{R}$ with $c_n - c_p \geq \delta$ and $c_p \leq t \leq c_n$ such that: (1) $b(t') \models_{\mathbb{R}} \pi$ for all $t' \in (c_p, c_n)$; (2) $b(c_n) \models_{\mathbb{R}} \square_{(0, \delta)}(\neg\pi) \vee \square(\pi)$; and (3) $b(c_p) \models_{\mathbb{R}} \overleftarrow{\square}_{(0, \delta)}(\neg\pi) \vee \overleftarrow{\square}(\pi)$. If in particular $c_n - c_p = \delta$ then also $b(c_n) \models_{\mathbb{R}} \pi$ and $b(c_p) \models_{\mathbb{R}} \pi$.*

PROOF. The proof follows easily from Definition 3, which entails that non-Berkeley behaviors $b \in \mathcal{BPR}_\delta$ are piecewise-constant functions of time whose discontinuities are at least δ time units apart. \square

2.3 Relations with Other Metric Temporal Logics

This section discusses expressiveness, decidability, and complexity results about MTL as has been introduced above.

2.3.1 Expressiveness. When defining the semantics of MTL formulas, several different choices are possible.

2.3.1.1 Global vs. initial satisfiability. First of all, notice that initial satisfiability is unambiguous only for mono-infinite time domains [Perrin and Pin 2004]. For such domains, it is clear that the global satisfiability semantics can be reduced to local satisfiability, as $b \models_{\mathbb{T}} \phi$ holds if and only if $b(0) \models_{\mathbb{T}} \text{Alw}(\phi)$ does. Conversely, local satisfiability is also reducible to global satisfiability, as for instance $b \models_{\mathbb{R}_{\geq 0}} \square_{>0}(\perp) \Rightarrow \phi$ is equivalent to $b(0) \models_{\mathbb{R}_{\geq 0}} \phi$ for the mono-infinite time domain $\mathbb{R}_{\geq 0}$, where $\square_{>0}(\perp)$ holds only at time 0. Therefore the two definitions of satisfiability are essentially equivalent for generic MTL formulas.

However, global satisfiability is significantly more expressive than initial satisfiability for *flat* \flat MTL formulas [Furia and Rossi 2007]. In particular, the expressiveness of the flat fragment is non-trivial under such global semantics as it corresponds to an implicit nesting of a qualitative temporal operator over the simpler initial satisfiability semantics. This entails that most common (real-time) properties — such as (bounded) response and (bounded) invariance [Koymans 1990] — can be easily expressed with flat formulas under the global satisfiability semantics. This is the main reason for adopting such a semantics in this paper whose results are focused on the flat fragment of MTL.

2.3.1.2 Flat vs. nesting. The syntactic restriction of flatness is also a semantic restriction, i.e., \flat MTL is strictly less expressive than full MTL. This is the case not only for dense time (which has been proved in [Furia and Rossi 2007]) but also for discrete time (which has been proved in [Etessami and Wilke 1996; Thérien and Wilke 2004; Kučera and Strejček 2005; Demri and Schnoebelen 2002] already for qualitative temporal logic), and regardless of whether a global or initial satisfiability semantics is assumed.

On the other hand, if we consider the weaker requirement of inter-reducibility of the satisfiability problems over global satisfiability, \flat MTL is as powerful as full MTL. In other words, given any MTL formula ϕ , it is possible to build a flat formula $\phi' \in \flat$ MTL which is globally satisfiable if and only if ϕ is. In general, ϕ' “flattens” ϕ by introducing additional propositional letters that are equivalent to matching nested sub-formulas in ϕ , as shown in the following.

EXAMPLE 9. Let $\phi = \mathbf{p} \Rightarrow \diamond_{<3}(\bigcirc(\square_{=2}(\mathbf{q})))$. Let us introduce the auxiliary propositions \mathbf{a}_1 and \mathbf{a}_2 defined as equivalent to $\square_{=2}(\mathbf{q})$ and $\bigcirc(\mathbf{a}_1)$ respectively. Hence, the derived flat formula

$$\phi' = (\mathbf{p} \Rightarrow \diamond_{<3}(\mathbf{a}_2)) \wedge (\mathbf{a}_1 \Leftrightarrow \square_{=2}(\mathbf{q})) \wedge (\mathbf{a}_2 \Leftrightarrow \bigcirc(\mathbf{a}_1))$$

is equi-satisfiable to ϕ under the global satisfiability semantics.

Details of this straightforward idea are shown in [Furia 2007; D’Souza et al. 2007] for dense time models, but it should be clear that a similar result can be proved for discrete time as well.

Let us finally consider dense-time behaviors that are non-Berkeley. In this case, the expressiveness gap between flat and nesting formulas still exists [Furia and Rossi 2007]. On the other hand, “flattening” is more intricate and cannot be done as with generic behaviors without breaking non-Berkeleyness as shown in the following example and discussed at greater length in Section 3.4.

EXAMPLE 10. *MTL formula $\psi = \text{Som}(\overleftarrow{\text{O}}(\neg p) \wedge \text{O}(p))$ describes behaviors where there exists a transition of proposition p from false to true. ψ is satisfiable over non-Berkeley behaviors \mathcal{BPR}_δ for any positive δ . However, consider the flattening $\bar{\psi}$ of ψ built according to the procedure described above.*

$$\bar{\psi} = \text{Som}(a) \wedge \left(a \Leftrightarrow \overleftarrow{\text{O}}(\neg p) \wedge \text{O}(p) \right)$$

Any behavior $b \in \mathcal{BPR}$ such that $b \models_{\mathbb{R}} \bar{\psi}$ requires a to hold at some instant t . However, sub-formula $a \Leftrightarrow \overleftarrow{\text{O}}(\neg p) \wedge \text{O}(p)$ forces a to hold only exactly at the transition points of p , pointwisely: any such b is Berkeley because $\overleftarrow{\text{O}}(\neg p) \wedge \text{O}(p)$ holds only at isolated points. Hence, ψ and $\bar{\psi}$ are not equi-satisfiable over non-Berkeley behaviors for any $\delta > 0$.

2.3.1.3 *Strictness and matchingness.* The semantics of an *until* formula with arguments ϕ_1, ϕ_2 requires the first argument ϕ_1 to hold over an interval $J = \langle 0, d \rangle \oplus t$ from current instant t . J can be taken to be open, half-open (with the left or right end-point included), or closed. Correspondingly four variants of *until* are possible. Each of them is labeled *strict* if J is open to the left and *non-strict* otherwise; and *matching* if J is closed to the right and *non-matching* otherwise [Furia and Rossi 2007]. The most common variant of the *until* operator is strict and non-matching, as it is simple to see that the three other variants are reducible to it. On the contrary, this paper adopts a non-strict non-matching *until* as basic operator, as the presentation of the results is more natural with non-strict operators.

In related work [Furia and Rossi 2007], we proved that all variants carry the same expressive power for MTL over dense- and discrete-time behaviors. On the contrary, strict *until* is more expressive than non-strict *until* for *flat* bMTL formulas.⁵

2.3.2 *Decidability and Complexity.* It is well-known that full MTL is undecidable over (non-Zeno) dense-time behaviors [Alur and Henzinger 1993]. The same holds for flat bMTL as its satisfiability problem is inter-reducible to the same problem for full MTL.

On the contrary, MTL becomes fully decidable over discrete time, with **EXPSpace**-complete complexity [Alur and Henzinger 1993]. MTL is also fully decidable over non-Berkeley dense-time behaviors \mathcal{BPT}_δ for any fixed δ , with the same complexity as over discrete time [Furia and Rossi 2008].

⁵[Furia and Rossi 2007] proves this for dense-time behaviors, but the same can be seen to hold over discrete-time behaviors as well.

3. SAMPLING INVARIANCE

Throughout this section we assume \mathbb{R} as dense (and continuous) time domain, and \mathbb{Z} as discrete time domain. It should be noted, however, that nearly all definitions and results can be adapted with little effort to work with different pairs of dense and discrete time domains as well, most notably the nonnegative reals and the naturals.

3.1 Definitions

3.1.1 Sampling Functions. A *sampling function* is a mapping between dense-time behaviors and discrete-time behaviors such that the latter are obtained by “sampling” — in some sense — the values of the former. We use $\zeta_{z,\delta}$ to denote a generic sampling function that is parametric with respect to a sampling period δ and an origin z .

The *canonical sampling* is a particular sampling function that models an idealized sampling process where a discrete-time behavior is obtained from a dense-time behavior by observing it at all instants corresponding to integer multiples of a chosen period δ .

DEFINITION 11 CANONICAL SAMPLING OF A BEHAVIOR. *Let $b \in \mathcal{BPR}$ be a dense-time behavior, $\delta \in \mathbb{R}_{>0}$ a positive real, and $z \in \mathbb{R}$ a basic offset. The canonical sampling $\sigma_{\delta,z}[b]$ of b is the discrete-time behavior in \mathcal{BPZ} defined by:*

$$\forall k \in \mathbb{Z} : \sigma_{\delta,z}[b](k) = b(z + k\delta)$$

We call δ the sampling period and z the origin of the sampling. Note that $\sigma_{\delta,z}$ is onto and total,⁶ for any δ, z .

Conversely, given a discrete-time behavior $d \in \mathcal{BPZ}$, $\sigma_{\delta,z}^{-1}[z]$ is the set of all dense-time behaviors such that their sampling is d .

$$\sigma_{\delta,z}^{-1}[d] = \{b \in \mathcal{BPR} \mid d = \sigma_{\delta,z}[b]\}$$

3.1.2 On Dense- vs. Discrete-Time Semantics. Consider some MTL formula ϕ that can be interpreted over both dense- and discrete-time behaviors. Its semantics is characterized by its dense-time models $\llbracket \phi \rrbracket_{\mathbb{R}}$ on the one hand, and by its discrete-time models $\llbracket \phi \rrbracket_{\mathbb{Z}}$ on the other hand. These two sets correspond to two different semantics for the *same* formula. The fact that the discrete time domain is a subset of the dense time domain prompts us to investigate the existence of a general relation linking the two sets $\llbracket \phi \rrbracket_{\mathbb{R}}$ and $\llbracket \phi \rrbracket_{\mathbb{Z}}$. More precisely, we seek simple conditions under which elements in $\llbracket \phi \rrbracket_{\mathbb{Z}}$ are precisely those obtained from elements in $\llbracket \phi \rrbracket_{\mathbb{R}}$ by applying the sampling function $\sigma_{\delta,z}$ for some δ, z .

This ideal requirement must be relaxed to some extent to be achievable in practice, for a number of reasons that are outlined informally in the following example.

EXAMPLE 12. *There are three fundamental discrepancies between discrete- and dense-time semantics that must be accommodated to reconcile them according to the notion of sampling.*

The first has to do with differences in terms of time units. Consider for instance formula $\Box_{\leq 2}(\mathbf{p})$; when interpreted over dense time, it states that \mathbf{p} holds for 2 time

⁶That is, it is defined for every $b \in \mathcal{BPR}$.

units. If we switch to a discrete-time interpretation and consider a sampling period of, say, $\delta = 3/10$, we would like the formula to refer to the same “sampled” interval. Hence, it should be changed to $\square_{\leq 20/3}(\mathbf{p})$ because the dense-time interval of length 2 becomes a discrete-time interval containing $2/(3/10)$ sampling instants.

However, $\square_{\leq 20/3}(\mathbf{p})$ cannot yet be interpreted over discrete time, as $20/3$ is not an integer; this shows a discrepancy in terms of granularity between dense and discrete sets. Of course, this problem can be solved by rounding the rational value to the nearest integer value, by taking its floor 6 or its ceiling 7. More precisely, whether to round up or down is decided in order to have a conservative approximation of the semantics. Intuitively, this means that intervals in “universal” formulas such as $\square_I(\mathbf{p})$ are rounded down — thus shrinking the interval into a smaller one —, whereas intervals in “existential” formulas such as $\diamond_I(\mathbf{p})$ are rounded up — thus expanding the interval into a larger one.

A similar granularity problem arises when interpreting a discrete-endpoint formula over dense time. Consider the example of formula $\diamond_{[1,2]}(\mathbf{p})$ that requires \mathbf{p} to hold one or two (discrete) time units in the future. In terms of dense-time units, \mathbf{p} must occur over the interval $[\delta, 2\delta] = [3/10, 3/5]$. However, the formula must hold also in between sampling instants when interpreted over dense-time behaviors. We will show that this feature of the dense-time semantics can be accommodated by expanding symmetrically the scaled interval into $[(1-1)\delta, (2+1)\delta] = [0, 9/10]$.

The last subtlety has to do with the change speed of dense-time behaviors with respect to the sampling period. Consider behavior b over proposition \mathbf{p} such that \mathbf{p} holds for less than δ time units, say over $[\delta/4, \delta/2]$. Formula $\text{Som}(\mathbf{p})$ is clearly satisfied by b over discrete time. However, any sampling $\sigma_{\delta,z}[b]$, for any $z \in (-\delta/2, \delta/4) \cup (\delta/2, \infty)$, does not have any sampling instant within $[\delta/4, \delta/2]$, and formula $\text{Som}(\mathbf{p})$ is not satisfied by any such $\sigma_{\delta,z}[b]$. This shows that only dense-time behaviors where state changes are sufficiently sparse can guarantee that formula satisfaction is preserved while moving to a sampled discrete-time semantics.

The discrepancies outlined above are bridged by introducing suitable notions. The concept of slowly-changing behavior is captured by the non-Berkeleyness constraint, introduced in Section 2.1. The following notion of *adaptation function* formalizes instead changes to intervals in MTL formulas, which take discrepancies between time units and granularities into account.

DEFINITION 13 ADAPTATION. A \mathbb{R} -to- \mathbb{Z} adaptation is a mapping from dense-endpoint to discrete-endpoint MTL formulas; a \mathbb{Z} -to- \mathbb{R} adaptation is a mapping from discrete-endpoint to dense-endpoint MTL formulas.

3.1.3 Sampling Invariance. We can finally introduce the definition of *sampling invariance* over non-Berkeley behaviors, which captures appropriately a notion of equivalence under sampling of models of MTL formulas.

DEFINITION 14 SAMPLING INVARIANCE. Let ϕ be an MTL formula over alphabet \mathcal{P} ; $v^{\mathbb{R}}, v^{\mathbb{Z}}$ a \mathbb{R} -to- \mathbb{Z} and \mathbb{Z} -to- \mathbb{R} adaptation, respectively; δ a sampling period; and $\varsigma_{\delta,z}$ a sampling function.

— ϕ is closed under sampling (c.u.s.) iff for any non-Berkeley behavior $b \in \mathcal{BPR}_{\delta}$

and any origin z :

$$b \in \llbracket \phi \rrbracket_{\mathbb{R}}^{\delta} \quad \text{implies} \quad \varsigma_{\delta,z}[b] \in \llbracket v^{\mathbb{R}}[\phi] \rrbracket_{\mathbb{Z}}$$

— ϕ is closed under inverse sampling (c.u.i.s.) iff for any discrete-time behavior $b \in \mathcal{BPZ}$ and any origin z :

$$b \in \llbracket \phi \rrbracket_{\mathbb{Z}} \quad \text{implies} \quad \forall b' \in \varsigma_{\delta,z}^{-1}[b] \cap \mathcal{BPR}_{\delta} \text{ it is } b' \in \llbracket v^{\mathbb{Z}}[\phi] \rrbracket_{\mathbb{R}}^{\delta}$$

— ϕ is sampling invariant (s.i.) iff it is c.u.s. if it is a dense-endpoint formula and it is c.u.i.s. if it is a discrete-endpoint formula.

Definition 14 depends on several parameters: \mathbb{R} -to- \mathbb{Z} and \mathbb{Z} -to- \mathbb{R} adaptations, a sampling period δ , and a sampling function $\varsigma_{\delta,z}$. In the following we will use the expression “sampling invariance (c.u.s. or c.u.i.s) with respect to” to highlight a particular choice for the parameters (when they are not obvious from the context).

3.2 Illustrative Examples

Before delving into the technical details of sampling invariance for generic MTL formulas, this sub-section illustrates the fundamental ideas that underlie the results of the paper. The presentation is deliberately partly informal and based on examples, with the goal of stimulating the intuition that substantiates the choice of adaptations (in Section 3.3.1) and the rationale of the technical proofs (in Section 15). In all the examples of this sub-section, we assume a sampling period $\delta = 1$.

The first example demonstrates the need for non-Berkeley behaviors with the same δ as the chosen sampling period. Consider formula $\diamond_{[2,5]}(\mathbf{p})$ and the behavior for \mathbf{p} in Figure 2(a). $\diamond_{[2,5]}(\mathbf{p})$ holds everywhere in dense time, but \mathbf{p} keeps on switching truth value in such a way that it is false at every sampled instant. If the sampling period is not commensurate to the “speed” of the dense-time behavior there is always the possibility of similarly twisted behaviors which prevent achieving c.u.s. even for very simple formulas. This justifies using the same δ for the non-Berkeley behaviors \mathcal{BPR}_{δ} and the sampling function $\varsigma_{\delta,z}$.

If we assume such a constraint on the behaviors considered, c.u.s. is straightforward for “existential” — that is “eventually” — formulas. Consider again formula $\diamond_{[2,5]}(\mathbf{p})$ and the behavior for \mathbf{p} in Figure 2(b). It should be clear that $c_2 \models \diamond_{[2,5]}(\mathbf{p})$ because \mathbf{p} holds at least once in any closed interval of length 3. It follows that the same holds for the discrete-time sampling d_2 of c_2 . In fact, consider any interval I of size 3 with integer endpoints and an instant within I where \mathbf{p} holds. Non-Berkeleyness entails that \mathbf{p} holds until the next sampling instant, since the previous sampling instant, or both. Hence, it reaches a sampling instant that fits the interval I over discrete time, which satisfies formula $\diamond_{[2,5]}(\mathbf{p})$ over discrete time. This can be generalized to show that no change in the time interval is required for existential formulas when passing from dense- to discrete-time interpretations — except for scaling the units according to the sampling period. As a concrete example in Figure 2(b), evaluate $\diamond_{[2,5]}(\mathbf{p})$ at -1 , which references the dense-time interval $[1, 4]$. Consider the instant between 1 and 2 marked with a cross where \mathbf{p} holds; \mathbf{p} also holds since 1, which is a sampling instant that belongs to the discrete-time interval $[1, 4]$.




Fig. 2. In all the pictures, the behavior of p is pictured by solid lines (in dense time) and discs (in sampled discrete time); the behavior of q is pictured by dotted lines (in dense time) and circles (in sampled discrete time); the higher value in any behavior corresponds to a \top truth value; for $i = 1, \dots, 6$, c_i denotes the dense-time behavior and d_i its discrete-time sampling.

A similar reasoning works for “universal” — that is “always” — formulas, such as $q \Rightarrow \Box_{[2,5]}(p)$. Behavior c_3 in Figure 2(c) is such that $c_3 \models q \Rightarrow \Box_{[2,5]}(p)$; in particular p has to hold over the dense-time interval $[1, 4]$. Over discrete-time, sampled values of p have to hold over the *discrete*-time interval with the same endpoints, which is obviously the case. Again, this generalizes to universal formulas, which do not require changes in the time intervals when adapting them from dense- to discrete-time interpretations.

Things are more convoluted for c.u.i.s., which mandates changing the size of the intervals according to the type of formula — existential or universal. Let us consider again the existential formula $\Diamond_{[2,5]}(p)$; it holds everywhere in the discrete-time behavior d_4 in Figure 2(d). If, however, the same formula is interpreted over the dense-time behavior c_4 , of which d_4 is a sampling, it does not hold everywhere. In particular, it holds at -2 and -1 but it does not hold in the open interval $(-2, -1)$: see the cross mark and the corresponding interval of size 3 starting between 0 and 1. The problem here is that non-Berkeleyness is a constraint on speed, not synchronization: the two samplings of p at 0 and 4 record the value of the dense-time behavior c_4 respectively right before 0 and right after 4, hence leaving it unconstrained in the open interval $(0, 4)$ of size larger than 3. The “interval of uncertainty” is never larger than one sampling period on each side, hence we suggest to introduce an \mathbb{Z} -to- \mathbb{R} adaptation that grows intervals in existential formulas by this amount, thus accommodating the uncertainty in the worst case. In the example, the adapted formula is $\Diamond_{[1,6]}(p)$ which clearly holds everywhere over c_4 .

The dual reasoning suggests the adaptation for universal formulas such as $q \Rightarrow \Box_{[2,5]}(p)$. In Figure 2(e), the formula holds everywhere over discrete time. Over dense time, however, q holds shortly before -1 (see cross mark) but p does not hold everywhere in the corresponding interval of size 3 starting shortly before 1. Again, a weaker formula holds over dense time, obtained by shrinking intervals in universal formulas by one sampling period on each side; the \mathbb{Z} -to- \mathbb{R} adaptation has to implement such a modification. In the example, the adapted formula is $q \Rightarrow \Box_{[3,4]}(p)$ which clearly holds everywhere over c_5 .

In order to rigorously extend the informal reasoning so far to arbitrary flat MTL formulas, we have to combine “eventually” and “always” formulas with the binary *until* and *release* modalities. Let us demonstrate the intuition behind handling the former which turns out to be more intricate. Consider a qualitative formula $U(q, p)$ and the behavior in Figure 2(f); let x denote the time instant between 1 and 2 marked with a cross and assume that p holds, in particular, precisely at x . Then, the *until* formula $U(q, p)$ holds continuously over the interval $(-\infty, 1]$ in dense time (and beyond up to x). Correspondingly, the same formula holds over the discrete interval $(-\infty, 1]$ in discrete time. This suggests that *until* formulas are c.u.s., as we will demonstrate formally in the rest of the paper.

Closure under inverse sampling is, again, more problematic. Consider the same formula $U(q, p)$ and the same discrete-time behavior d_6 ; we have seen that the *until* formula holds over the discrete interval $(-\infty, 1]$ in discrete time. Take a slightly different dense-time behavior, one where p is false and q is true at x and everything else is as in c_6 ; let us name c'_6 this modified behavior. Obviously, d_6 is a sampling of c'_6 as well as c_6 . However, $U(q, p)$ does not hold anywhere in $(-\infty, 1]$ over c'_6

because \mathbf{p} becomes true left-continuously at x , which is incompatible with the dense-time semantics of *until*. In this case, the \mathbb{Z} -to- \mathbb{R} adaptation will have to replace the second argument \mathbf{p} of the *until* formula with the weaker $\Delta(\mathbf{p})$ which holds at x in c'_6 (as well as in c_6). Alternatively, no adaptation is needed if we consider the stronger *matching* variant of *until* $\mathbf{U}^\downarrow(\mathbf{q}, \mathbf{p})$, where \mathbf{p} and \mathbf{q} would have to hold together at 1 or 2 in discrete time.

The following sub-sections present rigorous proofs of s.i. of MTL formulas that build upon the intuition behind the examples in the present sub-section.

3.3 Sampling Invariance for MTL

This section provides a proof of the following fundamental result: there exist two suitable regular adaptations $\eta_\delta^{\mathbb{R}}, \eta_\delta^{\mathbb{Z}}$ such that bMTL is s.i. for the canonical sampling $\sigma_{\delta, z}$. In addition, the adaptations can be proved to introduce *minimal* changes in the intervals of the adapted formulas, in the sense of Theorems 17 and 18 below.

3.3.1 Canonical Adaptations. Consider \mathbb{R} -to- \mathbb{Z} adaptation $\eta_\delta^{\mathbb{R}}$, parametric with respect to positive real parameter δ , defined inductively as follows.

$$\begin{aligned}
\eta_\delta^{\mathbb{R}}[\pi] &\triangleq \pi \\
\eta_\delta^{\mathbb{R}}\left[\mathbf{U}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{U}_{[\lfloor l/\delta \rfloor, \lceil u/\delta \rceil]}(\eta_\delta^{\mathbb{R}}[\phi_1], \eta_\delta^{\mathbb{R}}[\phi_2]) \\
\eta_\delta^{\mathbb{R}}\left[\mathbf{S}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{S}_{[\lfloor l/\delta \rfloor, \lceil u/\delta \rceil]}(\eta_\delta^{\mathbb{R}}[\phi_1], \eta_\delta^{\mathbb{R}}[\phi_2]) \\
\eta_\delta^{\mathbb{R}}\left[\mathbf{R}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{R}_{\langle l', u' \rangle}(\eta_\delta^{\mathbb{R}}[\phi_1], \eta_\delta^{\mathbb{R}}[\phi_2]) \\
&\text{where } l' = \begin{cases} \lfloor l/\delta \rfloor & \text{if } \langle \text{ is } (\\ \lceil l/\delta \rceil & \text{if } \langle \text{ is } [\\ \text{and } u' = \begin{cases} \lceil u/\delta \rceil & \text{if } \rangle \text{ is }) \\ \lfloor u/\delta \rfloor & \text{if } \rangle \text{ is }] \end{cases} \\
\eta_\delta^{\mathbb{R}}\left[\mathbf{T}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{T}_{\langle l', u' \rangle}(\eta_\delta^{\mathbb{R}}[\phi_1], \eta_\delta^{\mathbb{R}}[\phi_2]) \\
&\text{where } l' = \begin{cases} \lfloor l/\delta \rfloor & \text{if } \langle \text{ is } (\\ \lceil l/\delta \rceil & \text{if } \langle \text{ is } [\\ \text{and } u' = \begin{cases} \lceil u/\delta \rceil & \text{if } \rangle \text{ is }) \\ \lfloor u/\delta \rfloor & \text{if } \rangle \text{ is }] \end{cases} \\
\eta_\delta^{\mathbb{R}}[\phi_1 \wedge \phi_2] &\triangleq \eta_\delta^{\mathbb{R}}[\phi_1] \wedge \eta_\delta^{\mathbb{R}}[\phi_2] \\
\eta_\delta^{\mathbb{R}}[\phi_1 \vee \phi_2] &\triangleq \eta_\delta^{\mathbb{R}}[\phi_1] \vee \eta_\delta^{\mathbb{R}}[\phi_2]
\end{aligned}$$

Consider \mathbb{Z} -to- \mathbb{R} adaptation $\eta_\delta^{\mathbb{Z}}$, parametric with respect to positive real parameter δ , defined inductively as follows.⁷

⁷The restriction to closed intervals is clearly without loss of generality over discrete time.

$$\begin{aligned}
\eta_\delta^Z[\pi] &\triangleq \pi \\
\eta_\delta^Z\left[\mathbf{U}_{[l,u]}(\phi_1, \phi_2)\right] &\triangleq \mathbf{U}_{((l-2)\delta, (u+1)\delta)}(\eta_\delta^Z[\phi_1], \Delta(\eta_\delta^Z[\phi_2])) \\
\eta_\delta^Z\left[\mathbf{S}_{[l,u]}(\phi_1, \phi_2)\right] &\triangleq \mathbf{S}_{((l-2)\delta, (u+1)\delta)}(\eta_\delta^Z[\phi_1], \overleftarrow{\Delta}(\eta_\delta^Z[\phi_2])) \\
\eta_\delta^Z\left[\mathbf{R}_{[l,u]}(\phi_1, \phi_2)\right] &\triangleq \mathbf{R}_{[(l+1)\delta, (u-1)\delta]}(\eta_\delta^Z[\phi_1], \eta_\delta^Z[\phi_2]) \\
\eta_\delta^Z\left[\mathbf{T}_{[l,u]}(\phi_1, \phi_2)\right] &\triangleq \mathbf{T}_{[(l+1)\delta, (u-1)\delta]}(\eta_\delta^Z[\phi_1], \eta_\delta^Z[\phi_2]) \\
\eta_\delta^Z[\phi_1 \wedge \phi_2] &\triangleq \eta_\delta^Z[\phi_1] \wedge \eta_\delta^Z[\phi_2] \\
\eta_\delta^Z[\phi_1 \vee \phi_2] &\triangleq \eta_\delta^Z[\phi_1] \vee \eta_\delta^Z[\phi_2]
\end{aligned}$$

The proof of Theorem 15 will show that the asymmetry in the adaptation for *until* (and *since*) operators is needed to reconcile the non-matching semantics over discrete and dense time. Alternatively, one can assume discrete-endpoint *until* (and *since*) operators in their matching variant (see Table I) which preserves the symmetry in the adaptations. We include them explicitly in the treatment also because they will be useful for the results of Section 4.

$$\begin{aligned}
\eta_\delta^Z\left[\mathbf{U}_{[l,u]}^\downarrow(\phi_1, \phi_2)\right] &\triangleq \mathbf{U}_{((l-1)\delta, (u+1)\delta)}^\downarrow(\eta_\delta^Z[\phi_1], \eta_\delta^Z[\phi_2]) \\
\eta_\delta^Z\left[\mathbf{S}_{[l,u]}^\downarrow(\phi_1, \phi_2)\right] &\triangleq \mathbf{S}_{((l-1)\delta, (u+1)\delta)}^\downarrow(\eta_\delta^Z[\phi_1], \eta_\delta^Z[\phi_2])
\end{aligned}$$

We name η_δ^R and η_δ^Z *canonical* adaptations.

3.3.2 Flat MTL is Sampling Invariant. The main result of the paper is now proved.

THEOREM 15 SAMPLING INVARIANCE OF bMTL. *Let $\delta > 0$ be any sampling period and z be any origin. All flat bMTL formulas are sampling invariant with respect to the canonical adaptations $\eta_\delta^R, \eta_\delta^Z$ and the canonical sampling function $\sigma_{\delta,z}$.*

PROOF. The proof is split into two parts: first we show that any dense-endpoint flat formula ϕ is c.u.s.; then we show that any discrete-endpoint flat formula ϕ is c.u.i.s.⁸

Let us introduce the following abbreviations: for a dense-time instant r , let $\Omega(r)$ denote the sampling instant $z + \lfloor (r-z)/\delta \rfloor \delta$, which is immediately before or exactly at r , and let $O(r)$ denote the sampling instant $z + \lceil (r-z)/\delta \rceil \delta$, which is immediately after or exactly at r . Also, $\omega(r)$ and $o(r)$ denote the distances between r and its previous and next sampling instant, respectively; that is $\omega(r) = r - \Omega(r)$ and $o(r) = O(r) - r$. Obviously $\omega(r), o(r) \geq 0$.⁹

(*Closure under sampling*). Let ϕ be a generic dense-endpoint flat MTL formula, b a dense-time non-Berkeley behavior in $[[\phi]]_{\mathbb{R}}^\delta$, and $\phi' = \eta_\delta^R[\phi]$. Then, let b' be the sampling $\sigma_{\delta,z}[b]$ of b with the given origin and sampling period.

⁸For brevity we omit dealing with past operators, as it can be done from the corresponding future operators with little effort.

⁹This proof exploits some properties of the floor and ceiling functions. We refer the reader to [Graham et al. 1994] for a thorough treatment of these functions.

For a generic sampling instant $t = z + k\delta$, we show that $b(t) \models_{\mathbb{R}} \phi$ implies $b'(k) \models_{\mathbb{Z}} \phi'$, by induction on the structure of ϕ . This proves that if $b \models_{\mathbb{R}} \phi$ then $\sigma_{\delta, z}[b] \models_{\mathbb{Z}} \eta_{\delta}^{\mathbb{R}}[\phi]$, for any $b \in \mathcal{BPR}_{\delta}$; hence any bMTL dense-endpoint formula is c.u.s.

— $\phi = \pi$, $\phi = \phi_1 \wedge \phi_2$, and $\phi = \phi_1 \vee \phi_2$ are straightforward from the definitions.

— $\phi = \mathbf{U}_{\langle l, u \rangle}(\pi_1, \pi_2)$. ϕ' is $\mathbf{U}_{[l', u']}(\pi_1, \pi_2)$, with $l' = \lfloor l/\delta \rfloor$ and $u' = \lceil u/\delta \rceil$.

Let d be a real in $\langle l, u \rangle$ such that $b(t+d) \models_{\mathbb{R}} \pi_2$ and, for all $e \in [0, d)$, it is $b(t+e) \models_{\mathbb{R}} \pi_1$. Since b in non-Berkeley, there exists a $p \in [0, \delta]$ such that for all $f \in [-p, -p+\delta]$ it is $b(t+d+f) \models_{\mathbb{R}} \pi_2$; i.e., π_2 holds over $I = [t+d-p, t+d-p+\delta]$. Some sampling instant must fall within I , as I has size δ .

In particular, it is either $p \geq \omega(t+d)$ or $-p+\delta \geq \mathfrak{o}(t+d)$: otherwise it would be $\delta = p + (-p+\delta) < \omega(t+d) + \mathfrak{o}(t+d) = \mathbf{O}(t+d) - \mathbf{\Omega}(t+d) = \delta(\lceil (t+d-z)/\delta \rceil - \lfloor (t+d-z)/\delta \rfloor) \leq \delta$, a contradiction (where we exploited the property: $\lceil r \rceil - \lfloor r \rfloor \leq 1$ for any real r). So, let t' be the sampling instant:

$$t' = \begin{cases} \mathbf{\Omega}(t+d) & \text{if } p \geq \omega(t+d) \\ \mathbf{O}(t+d) & \text{otherwise} \end{cases}$$

It is not difficult to check that $(t' - t)/\delta \in [l', u']$. In fact:

—if $p \geq \omega(t+d)$, then $t' - t = \mathbf{\Omega}(t+d) - t = \delta(\lfloor (k\delta+d)/\delta \rfloor - k) = \delta\lfloor d/\delta \rfloor$. Recall that $d \in \langle l, u \rangle$, and then *a fortiori* $d \in [l, u] \supseteq \langle l, u \rangle$. So $d/\delta \in [l/\delta, u/\delta]$, and $(t' - t)/\delta = \lfloor d/\delta \rfloor \in [\lfloor l/\delta \rfloor, \lfloor u/\delta \rfloor] \subseteq [l', u']$.

—if $p < \omega(t+d)$, then $t' - t = \mathbf{O}(t+d) - t = \delta(\lceil (k\delta+d)/\delta \rceil - k) = \delta\lceil d/\delta \rceil$. Recall that $d \in \langle l, u \rangle$, and then *a fortiori* $d \in [l, u] \supseteq \langle l, u \rangle$. So $d/\delta \in [l/\delta, u/\delta]$, and $(t' - t)/\delta = \lceil d/\delta \rceil \in [\lceil l/\delta \rceil, \lceil u/\delta \rceil] \subseteq [l', u']$.

In all, $b(t') \models_{\mathbb{R}} \pi_2$. By inductive hypothesis, it follows that for $d' = (t' - t)/\delta$ it is $b'(k+d') \models_{\mathbb{Z}} \pi_2$, and $d' \in [l', u']$.

Let us now show that for all integers $e' \in [0, d' - 1]$ it is $b'(k+e') \models_{\mathbb{Z}} \pi_1$. Recall that $d' \leq \lceil d/\delta \rceil < d/\delta + 1$, since $\lceil r \rceil < r + 1$ for any real number r ; hence $\delta(d' - 1) < \delta(d/\delta) = d$. Since for all $e \in [0, d)$ we have $b(t+e) \models_{\mathbb{R}} \pi_1$, and since $[0, \delta(d' - 1)] \subset [0, d)$, *a fortiori* for all $e \in [0, \delta(d' - 1)]$ it is $b(t+e) \models_{\mathbb{R}} \pi_1$. By inductive hypothesis, it follows that for all integers $e' \in [0, d' - 1] = [0, d')$ it is $b'(k+e') \models_{\mathbb{Z}} \pi_1$. We conclude that $b'(k) \models_{\mathbb{Z}} \phi'$.

— $\phi = \mathbf{R}_{\langle l, u \rangle}(\pi_1, \pi_2)$. ϕ' is $\mathbf{R}_{[l', u']}(\pi_1, \pi_2)$, where l', u' depend on whether $I = \langle l, u \rangle$ is closed, open, or half-open.

Let d' be a generic integer in $\langle l', u' \rangle$; we show that $b'(k+d') \models_{\mathbb{Z}} \pi_2$ or there exists a $e' \in [0, d')$ such that $b'(k+e') \models_{\mathbb{Z}} \pi_1$. First we show that $\langle l', u' \rangle \subseteq [l/\delta, u/\delta]$. In fact, consider the four possible cases for interval $I' = \langle l', u' \rangle$.

— $I = [l, u]$, so $I' = [l', u']$, where $l' = \lceil l/\delta \rceil$ and $u' = \lfloor u/\delta \rfloor$. Thus, $[l', u'] \subseteq [l/\delta, u/\delta]$, as $\lfloor r \rfloor \leq r$ and $\lceil r \rceil \geq r$ for any real r .

— $I = [l, u)$, so $I' = [l', u')$, where $l' = \lceil l/\delta \rceil$ and $u' = \lfloor u/\delta \rfloor$. Thus, $[l', u') \subseteq [l/\delta, u/\delta)$, as $[l', u') = [\lceil l/\delta \rceil, \lfloor u/\delta \rfloor - 1] \subseteq [l/\delta, u/\delta)$, noting that $\lceil r \rceil \geq r$, and that $\lfloor r \rfloor - 1 < r$, for any real r .

— $I = (l, u]$, so $I' = (l', u']$, where $l' = \lfloor l/\delta \rfloor$ and $u' = \lceil u/\delta \rceil$. Thus, $(l', u'] \subseteq (l/\delta, u/\delta]$, as $(l', u'] = [\lfloor l/\delta \rfloor + 1, \lceil u/\delta \rceil] \subseteq (l/\delta, u/\delta]$, noting that $\lfloor r \rfloor \leq r$, and that $\lfloor r \rfloor + 1 > r$, for any real r .

— $I = (l, u)$, so $I' = (l', u')$, where $l' = \lfloor l/\delta \rfloor$ and $u' = \lceil u/\delta \rceil$. Thus, $(l', u') \subseteq (l/\delta, u/\delta)$, as $(l', u') = [\lfloor l/\delta \rfloor + 1, \lceil u/\delta \rceil - 1] \subset (l/\delta, u/\delta)$, noting that $\lfloor r \rfloor + 1 > r$, and that $\lceil r \rceil - 1 < r$, for any real r .

In all, $b(t + \delta d') \models_{\mathbb{R}} \pi_2$ or there exists a $e \in [0, \delta d')$ such that $b(t + e) \models_{\mathbb{R}} \pi_1$.

If the former is the case, $b'(k + d') \models_{\mathbb{Z}} \pi_2$ holds by inductive hypothesis, which fulfills the goal. If the latter is the case, we have $b(t) \models_{\mathbb{R}} \diamond_{[0, \delta d')}(\pi_1) \equiv \bigcup_{[0, \delta d')}(\top, \pi_1)$, which entails $b'(k) \models_{\mathbb{Z}} \diamond_{[0, d')}(\pi_1)$. Therefore, there exists a $e' \in [0, d')$ such that $b'(k + e') \models_{\mathbb{Z}} \pi_1$, as required.

(*Closure under inverse sampling*). Let us first introduce the following terminology; for any dense-endpoint formula ψ :

- if $b(t) \models_{\mathbb{R}} \square_{<\delta}(\psi)$ (resp. $b(t) \models_{\mathbb{R}} \overleftarrow{\square}_{<\delta}(\psi)$), ψ “shifts to the right (s.t.r.) at t ” (resp. “shifts to the left (s.t.l.) at t ”);
- if $b(t) \models_{\mathbb{R}} \mathbf{U}_{=c}(\psi, \Delta(\neg\psi))$ (resp. $b(t) \models_{\mathbb{R}} \mathbf{S}_{=c}(\psi, \overleftarrow{\Delta}(\neg\psi))$) for some $c \in (0, \delta)$, or $b(t) \models_{\mathbb{R}} \psi \wedge \bigcirc(\neg\psi)$ (resp. $b(t) \models_{\mathbb{R}} \psi \wedge \overleftarrow{\bigcirc}(\neg\psi)$) and $c = 0$, ψ “turns false in the future (t.f.f.) at $t \stackrel{\pm}{\rightarrow} c$ ” (resp. “turned false in the past (t.f.p.) at $t \overrightarrow{-} c$ ”).

Let ϕ be a generic discrete-endpoint flat MTL formula, b a discrete-time behavior in $\llbracket \phi \rrbracket_{\mathbb{Z}}$, and $\phi' = \eta_{\delta}^z[\phi]$. Then, let b' be a dense-time non-Berkeley behavior in \mathcal{BPR}_{δ} such that $\sigma_{\delta, z}[b'] = b$ with the given origin and sampling period.

For a generic sampling instant $t = z + k\delta$, in the remainder we show that: (1) $b'(t) \models_{\mathbb{R}} \phi'$; (2) ϕ' s.t.r. at t , or there exist $c \in [0, \delta)$ and $\varpi \in \text{PL}$ such that ϕ' and ϖ both t.f.f. at $t \stackrel{\pm}{\rightarrow} c$, or ϕ' is false at $t + \delta$; and (3) either ϕ' s.t.l. at t or there exist $c \in [0, \delta)$ and $\varpi \in \text{PL}$ such that ϕ' and ϖ both t.f.p. at $t \overrightarrow{-} c$.

From these three facts we can prove that ϕ is c.u.i.s. by showing that $b'(t) \models_{\mathbb{R}} \phi'$ for all $t \in \mathbb{R}$. First, (1) shows this fact for all $t = z + k\delta$ for some integer k . Then, let $t_n = t + \delta$ and show that ϕ' holds over the generic δ -length closed real interval $[t, t_n]$. If ϕ' s.t.r. at t or it s.t.l. at t_n , we are done. If ϕ' is false at $t + \delta = z + (k+1)\delta$ we have a contradiction which also closes the proof. Otherwise, from (2) and (3) we assume that: (a) ϕ' t.f.f. at $t \stackrel{\pm}{\rightarrow} c_p$ for some $c_p \in [0, \delta)$ with some $\varpi_p \in \text{PL}$; and (b) ϕ' t.f.p. at $t_n \overrightarrow{-} c_n$ for some $c_n \in -[0, \delta)$ with some $\varpi_n \in \text{PL}$. Note that $|(t_n + c_n) - (t + c_p)| = |\delta + c_n - c_p| \leq \delta$; non-Berkeyness of $b' \in \mathcal{BPR}_{\delta}$ entails that either the two change points $t + c_p$ and $t_n + c_n$ coincide or $c_n = c_p = 0$. In both cases Lemma 8 implies a contradiction which closes the whole proof. We remark that the proofs go through also for intervals of temporal operators with negative endpoints, possibly with minimal adjustments that we do not discuss explicitly for the sake of brevity.

Finally, we prove (1), (2), and (3) by induction on the structure of ϕ .

— $\phi = \pi$. (1) From the definition of $\sigma_{\delta, z}$, it follows that $b'(t) = b(k)$.

(2) Consider Lemma 8 at t : there exist $c_n \geq t$ such that π t.f.f. at $t \stackrel{\pm}{\rightarrow} c_n$ or it holds indefinitely in the future. If the latter is the case, π obviously s.t.r. (3) is proved similarly as (2), with respect to the past.

— $\phi = \phi_1 \wedge \phi_2$ and $\phi = \phi_1 \vee \phi_2$ are straightforward from the definitions.

- $-\phi = \mathbf{U}_{[l,u]}(\pi_1, \pi_2)$. ϕ' is $\mathbf{U}_{(l',u')}(\pi_1, \Delta(\pi_2))$, with $l' = (l-2)\delta$ and $u' = (u+1)\delta$.
- (1) Let us start by proving $b'(t) \models_{\mathbb{R}} \psi$ with $\psi = \mathbf{U}_{[(l-1)\delta, u\delta]}(\pi_1, \pi_2)$. This implies $b'(t) \models_{\mathbb{R}} \phi'$, as $[(l-1)\delta, u\delta] \subset (l', u')$ hence ψ is stronger than ϕ' . Proving a stronger formula will be necessary in steps (2) and (3).
- Let $d \in [l, u]$ be the integer time instant such that $b(k+d) \models_{\mathbb{Z}} \pi_2$, which exists by hypothesis. The case $d = 0$ is trivial, hence let us consider $d > 0$. Still by hypothesis, for all integers $e \in [0, d] = [0, d-1]$ it is $b(k+e) \models_{\mathbb{Z}} \pi_1$. By inductive hypothesis, for all real δ -multiples $e' \in [0, d\delta]$ it is $b'(t+e') \models_{\mathbb{R}} \pi_1$. If $b'(t+d\delta) \models_{\mathbb{R}} \pi_1$ as well, let $d' = d\delta$; otherwise π_1 t.f.f. at some $t' \stackrel{+}{\rightarrow} 0$ with $t + (d-1)\delta \leq t' \leq t + d\delta$ and let $d' = t' - t \geq 0$. Notice that $d' \in [(l-1)\delta, u\delta]$. Correspondingly, π_1 holds over $[0, d'] \oplus t$. In addition, a little reasoning should convince us that Lemma 8 for π_2 at $t+d\delta$ — also considering the fact that π_1 t.f.f. at $t' \stackrel{+}{\rightarrow} 0$ unless it holds at $t+d\delta$ — implies that π_2 must hold over $(d', d\delta] \oplus t$; hence $b'(t+d') \models_{\mathbb{R}} \Delta(\pi_2)$.
- (2) Let s be any value in $(0, \delta)$. Let $c = d' - s$: notice that $c \in (l', u')$ because $c > d' - \delta \geq (l-1)\delta - \delta = l'$ and $c < d' \leq u\delta < u'$. Since $t + s + c = t + d'$ we have already shown that $b'(t+s+c) \models_{\mathbb{R}} \Delta(\pi_2)$. Moreover, $[s, s+c] \subset [0, d']$ thus $b'(t+s+f) \models_{\mathbb{R}} \pi_1$ holds *a fortiori* for all $f \in [0, c]$. All this proves ϕ' s.t.r.
- (3) Let f be any value in $-(0, \delta)$. For $d'' = d' - f$ we have $d'' \in [(l-1)\delta, (u+1)\delta]$ and $b'(t+f+d'') \models_{\mathbb{R}} \Delta(\pi_2)$. Also, by inductive hypothesis either π_1 t.f.p. at $t \stackrel{-}{\rightarrow} c$ for some $c \in [0, \delta)$ or π_1 s.t.l. In the latter case, ϕ' s.t.l. as well; in the former case, ϕ' t.f.p. at $t \stackrel{-}{\rightarrow} c$ as well.
- $-\phi = \mathbf{R}_{[l,u]}(\pi_1, \pi_2)$. ϕ' is $\mathbf{R}_{[l',u']}(\pi_1, \pi_2)$, with $l' = (l+1)\delta$ and $u' = (u-1)\delta$.
- (1) Let us start by proving $b'(t) \models_{\mathbb{R}} \psi$ with $\psi = \mathbf{R}_{[l\delta, u\delta]}(\pi_1, \pi_2)$. This implies $b'(t) \models_{\mathbb{R}} \phi'$, as $[l\delta, u\delta] \supset [l', u']$ hence ψ is stronger than ϕ' . Proving a stronger formula will be necessary in steps (2) and (3).
- Let d' be any real value in $[l\delta, u\delta]$; we prove that $b'(t+d') \models_{\mathbb{R}} \pi_2$ or $b'(t+e') \models_{\mathbb{R}} \pi_1$ for some $e' \in [0, d')$. We discuss two cases.
- If $t+d'$ is a sampling instant, $d = d'/\delta$ is an integer, and $d \in [l, u]$. Also, by hypothesis, $b(k+d) \models_{\mathbb{Z}} \pi_2$ or $b(k+e) \models_{\mathbb{Z}} \pi_1$ for some integer $e \in [0, d-1]$. In the former case, $b'(t+d') \models_{\mathbb{R}} \pi_2$ follows by inductive hypothesis. Otherwise, $b'(t+e') \models_{\mathbb{R}} \pi_1$ for $e' = e\delta$ and $e' \in [0, d' - \delta] \subset [0, d')$, also by inductive hypothesis.
- If $t+d'$ is not a sampling instant, let $p' = d' - \omega(t+d')$ and $n' = d' + \omega(t+d')$; these are both integer multiples of δ . Notice that $p' > d' - \delta \geq l\delta - \delta = (l-1)\delta$, and $n' < d' + \delta \leq u\delta + \delta = (u+1)\delta$. Therefore, the two integers $p = p'/\delta$ and $n = n'/\delta$ are such that $p, n \in [l, u]$. Hence, from the hypothesis $b(k) \models_{\mathbb{R}} \phi$, one of the following two cases holds.
- $b(k+p) \models_{\mathbb{Z}} \pi_2$ and $b(k+n) \models_{\mathbb{Z}} \pi_2$, with $n = p+1$. By inductive hypothesis, $b'(t+p') \models_{\mathbb{R}} \pi_2$ and $b'(t+n') \models_{\mathbb{R}} \pi_2$ follow. Since $b' \in \mathcal{BPR}_\delta$ is non-Berkeley by hypothesis, π_2 holds over the whole real interval $[t+p', t+n'] = [t+p', t+p'+\delta]$ as well. In particular, $b'(t+d') \models_{\mathbb{R}} \pi_2$ for $d' \in [p', p'+\delta]$.
- $b(k+e) \models_{\mathbb{Z}} \pi_1$ for some integer $e \in [0, p-1]$ or $e \in [0, n-1]$. From $(p-1)\delta \leq (n-1)\delta < (d'+\delta) - \delta = d'$, it follows that $e' = e\delta \in [0, d')$. $b'(t+e') \models_{\mathbb{R}} \pi_1$ holds by inductive hypothesis.

In all, $b'(t) \models_{\mathbb{R}} \psi$ is established.

(2) Let f be any value in $(0, \delta)$ and d'' be any real value in $[l', u']$. Since $d'' + f \in [l\delta, u\delta]$, we already proved that $b'(t + d'' + f) \models_{\mathbb{R}} \pi_2$ or $b'(t + c) \models_{\mathbb{R}} \pi_1$ for some $c \in [0, d'' + f]$.

If, for all d'' , the *stronger* fact that $b'((t + f) + d'') \models_{\mathbb{R}} \pi_2$ or $b'(t + c) \models_{\mathbb{R}} \pi_1$ for some $c \in [f, d'' + f] \subset [0, d'' + f]$ holds, then we have proved that ϕ' s.t.r. at t — because $t + c = t + f + (c - f)$ and $c - f \in [0, d'']$.

Otherwise, there is some d'' such that: (a) $b'((t + f) + d'') \models_{\mathbb{R}} \neg\pi_2$; (b) $b'(t + c') = b'((t + f) + (c' - f)) \models_{\mathbb{R}} \neg\pi_1$ for all $c' \in [f, d'' + f]$; and (c) $b'(t + c) \models_{\mathbb{R}} \pi_1$ for some $c \in [0, f]$. Let v be the smallest instant in $[0, f)$ such that $\Delta(\neg\pi_1)$ holds at $t + v$; this exists because b' is non-Zeno. Lemma 8 entails that π_1 holds over interval $[0, v) \oplus t$, and π_1 t.f.f. at $t \xrightarrow{+} v$. Hence, it can be seen that ϕ' t.f.f. at $t \xrightarrow{+} v$ as well.

(3) Let s be any value in $-(0, \delta)$ and d be any real value in $[l', u']$. Since $s + d \in [l\delta, u\delta]$, we have already shown that $b'(t + (s + d)) \models_{\mathbb{R}} \pi_2$ or $b'(t + e') \models_{\mathbb{R}} \pi_1$ for some $e' \in [0, s + d)$. In both cases it follows that ϕ' s.t.l. at t , in particular as $e'' = e' - s$ with $e'' \in [-s, d) \subset [0, d)$ and $t + e' = (t + s) + e''$.

— $\phi = \mathbf{U}_{[l, u]}^{\downarrow}(\pi_1, \pi_2)$.

Proof is all similar to the case of the “standard” *until* with the simplification that matchingness allows us to establish the stronger $b'(t) \models_{\mathbb{R}} \mathbf{U}_{[l\delta, u\delta]}^{\downarrow}(\pi_1, \pi_2)$ in part (1). \square

3.3.3 Canonical Adaptations are Optimal. Let us provide some justification for the particular choice of canonical adaptations. In principle, more complex transformations could be devised such that Theorem 3.3 still holds. However, we aimed at introducing adaptations that change the structure of the formulas as little as possible, such that the transformed formulas are “essentially the same” as the original formulas, except for some adjustments required to bridge the gaps in terms of time units and granularity (see Example 12).

In a nutshell, adaptations should preserve the propositional and modal structure of a formula as much as possible. To formalize this intuition we introduce the notion of regularity.

DEFINITION 16 REGULARITY OF ADAPTATIONS. *Let ϕ_1, ϕ_2 any pair of MTL formulas and \mathbf{O} a modality. An adaptation v is:*

- Compositional *if it satisfies $v[\phi_1 \wr \phi_2] \equiv v[\phi_1] \wr v[\phi_2]$ for any $\wr \in \{\wedge, \vee\}$.*
- Propositional-preserving *if $v[\mathbf{p}] \equiv \mathbf{p}$ for any $\mathbf{p} \in \mathcal{P}$.*
- \mathbf{O} -modality-preserving *if, for any interval I , $v[\mathbf{O}_I(\phi_1, \phi_2)] \equiv \mathbf{O}_I(v[\phi_1], v[\phi_2])$.*

An adaptation is \mathbf{O} -regular if it is compositional, propositional-preserving, and \mathbf{O} -modality-preserving. An adaptation is regular when it is \mathbf{O} -regular for every modality $\mathbf{O} \in \{\mathbf{U}, \mathbf{S}, \mathbf{R}, \mathbf{T}\}$.

Canonical adaptations $\eta_{\delta}^{\mathbb{R}}, \eta_{\delta}^{\mathbb{Z}}$ are regular for all modalities, with the exception of $\eta_{\delta}^{\mathbb{Z}}$ which is not \mathbf{U} -modality-preserving for the non-matching variant of the *until* modality. A \mathbf{U} -modality-preserving \mathbb{Z} -to- \mathbb{R} adaptation, however, would not achieve sampling invariance: as noted in Section 3.3.1, the matching semantics is the most

natural choice to bridge the discrete- and dense-time semantics. Furthermore, canonical adaptations are the “best” among all possible regular sampling-invariant adaptations, in the sense that the adapted intervals are as constraining as possible. This should be intuitively understandable already from the proof of Theorem 15, which would not stand if we introduced any relaxation in adapted interval bounds. More formally we have the following.

THEOREM 17 OPTIMALITY OF $\eta_\delta^{\mathbb{R}}$. *Let $v^{\mathbb{R}}$ be a regular \mathbb{R} -to- \mathbb{Z} adaptation such that any flat dense-endpoint $\phi \in \text{bMTL}$ is c.u.s. with respect to it and $\sigma_{\delta,z}$. Then, $d \models_{\mathbb{Z}} \eta_\delta^{\mathbb{R}}[\phi]$ implies $d \models_{\mathbb{Z}} v^{\mathbb{R}}[\phi]$ for any behavior $d \in \mathcal{BPZ}$.*

PROOF PROOF. The proof relies on techniques very similar to those of Theorem 15, hence only a proof sketch is provided.

The proof goes by contradiction: let $v^{\mathbb{R}}$ be a \mathbb{R} -to- \mathbb{Z} regular adaptation such that there exist $\phi \in \text{bMTL}$ and $d \in \mathcal{BPZ}$ with $d \models_{\mathbb{Z}} \eta_\delta^{\mathbb{Z}}[\phi]$ but $d \not\models_{\mathbb{Z}} v^{\mathbb{R}}[\phi]$. Then, we build $c' \in \mathcal{BPR}_\delta$ and $\zeta \in \text{bMTL}$ such that $c' \models_{\mathbb{R}} \zeta$, $\sigma_{\delta,z}[c'] \models_{\mathbb{Z}} \eta_\delta^{\mathbb{R}}[\zeta]$, but $\sigma_{\delta,z}[c'] \not\models_{\mathbb{Z}} v^{\mathbb{R}}[\zeta]$; hence ζ is not c.u.s. with respect to $v^{\mathbb{R}}$ and $\sigma_{\delta,z}$.

Let $k \in \mathbb{Z}$ be such that $d(k) \not\models_{\mathbb{Z}} v^{\mathbb{R}}[\phi]$, while recall that $d(k) \models_{\mathbb{Z}} \eta_\delta^{\mathbb{R}}[\phi]$. The propositional structure of $\eta_\delta^{\mathbb{R}}[\phi]$ and $v^{\mathbb{R}}[\phi]$ is the same, since both adaptations are regular. Then, by induction on the same propositional structure of $\eta_\delta^{\mathbb{R}}[\phi]$ and $v^{\mathbb{R}}[\phi]$, one can show that there exists a *modality* $\mathbf{O} \in \{\mathbf{U}, \mathbf{R}, \mathbf{S}, \mathbf{T}\}$ such that $d(k) \models_{\mathbb{Z}} \eta_\delta^{\mathbb{R}}[\mathbf{O}_I(\pi_1, \pi_2)]$ and $d(k) \not\models_{\mathbb{Z}} v^{\mathbb{R}}[\mathbf{O}_I(\pi_1, \pi_2)]$ for some $\pi_1, \pi_2 \in \text{PL}$. Let us write $\mathbf{O}_J(\beta_1, \beta_2)$ for $\eta_\delta^{\mathbb{R}}[\mathbf{O}_I(\pi_1, \pi_2)]$, and $\mathbf{O}_K(\gamma_1, \gamma_2)$ for $v^{\mathbb{R}}[\mathbf{O}_I(\pi_1, \pi_2)]$. The proof goes on by case discussion on the modality \mathbf{O} ; for brevity we just show the case $\mathbf{O} = \mathbf{U}$, but the remaining cases can be handled all similarly.

Let i be one of 1, 2. From the definition of $\eta_\delta^{\mathbb{R}}$ and the regularity of $v^{\mathbb{R}}$, it is $\beta_i \equiv \gamma_i \equiv \pi_i$. In all, there exists a $u \in J$ s.t. $d(k+u) \models_{\mathbb{Z}} \pi_2$ and $d(h) \models_{\mathbb{Z}} \pi_1$ for all $h \in [0, u] \oplus k$. Since we are assuming that $d(k) \not\models_{\mathbb{Z}} \mathbf{U}_K(\gamma_1, \gamma_2)$, it must be $u \notin K$, so either $K \subseteq (-\infty, u-1]$ or $K \subseteq [u+1, +\infty)$. The remainder assumes $K \subseteq (-\infty, u-1]$ and $u > 1$; the other cases can be handled along the same lines and are omitted for brevity.

The next step builds a new formula $\varphi \triangleq \mathbf{U}_I(y, z)$ with fresh propositional letters y, z ; and a new discrete-time behavior e over $\{y, z\}$ defined as follows: $z \in e(j)$ iff $j \geq k+u$, and $y \in e(j)$ for all $j \in \mathbb{Z}$. It follows that $e(k) \models_{\mathbb{Z}} \mathbf{U}_J(y, z)$ but $e(k) \not\models_{\mathbb{Z}} \mathbf{U}_K(y, z)$ because $\max K < u$. Also notice that $\eta_\delta^{\mathbb{R}}[\varphi] = \mathbf{U}_J(y, z)$ and $v^{\mathbb{R}}[\varphi] = \mathbf{U}_K(y, z)$. Take a c built as follows: $z \in c(t)$ iff $t > z + (k+u-1)\delta$ and $y \in c(t)$ for all $t \in \mathbb{R}$. It should be clear that $c \in \mathcal{BPR}_\delta$, $c \in \sigma_{\delta,z}^{-1}[e]$, and $c(z+k\delta) \models_{\mathbb{R}} \mathbf{U}_I(y, z) = \varphi$, because z holds to the right of $z + (k+u-1)\delta$ but is false before and at it (over $z + \delta(k \oplus K)$). In addition, one can see that $c(t) \models_{\mathbb{R}} \varphi$ holds for all $t \geq z + k\delta$.

The last step is as follows: let us build a new non-Berkeley dense-time behavior $c' \in \mathcal{BPR}_\delta$ over propositions in $\{y, z\} \cup \{x\}$, i.e., the same propositions as c plus a fresh one denoted by x . $c'|_{\mathcal{P}}$ is identical to c , whereas $x \in c'(t)$ iff $c(t) \not\models_{\mathbb{R}} \varphi$; hence in particular $c'(z+k\delta) \models_{\mathbb{R}} \varphi \wedge \neg x$. Notice that such c' is non-Berkeley for δ . Finally, consider formula $\zeta \in \text{bMTL}$ defined as $x \vee \varphi$. Clearly, $c' \models_{\mathbb{R}} \zeta$ is the case by construction; hence $\sigma_{\delta,z}[c'] \models_{\mathbb{Z}} \eta_\delta^{\mathbb{R}}[\zeta]$ follows from Theorem 15. Also, $\sigma_{\delta,z}[c'](k) \not\models_{\mathbb{Z}} v^{\mathbb{R}}[\varphi]$ as the truth of φ does not depend on x ; and $\sigma_{\delta,z}[c'](k) \not\models_{\mathbb{Z}}$

$v^{\mathbb{R}}[x]$ as the regularity of $v^{\mathbb{R}}$ implies $v^{\mathbb{R}}[x] = x$ and $c'(z + k\delta) \not\models_{\mathbb{R}} x$. In all we have $c' \models_{\mathbb{R}} \zeta$, $\sigma_{\delta,z}[c'] \models_{\mathbb{Z}} \eta_{\delta}^{\mathbb{R}}[\zeta]$, and $\sigma_{\delta,z}[c'] \not\models_{\mathbb{Z}} v^{\mathbb{R}}[\zeta]$. Hence, c.u.s. does not hold for formula ζ with respect to adaptation v and $\sigma_{\delta,z}$, which is the desired contradiction. \square

With a very similar approach the following theorem about $\eta_{\delta}^{\mathbb{Z}}$ adaptation can be proved.

THEOREM 18 OPTIMALITY OF REGULAR $\eta_{\delta}^{\mathbb{Z}}$. *Let $v^{\mathbb{Z}}$ be a regular \mathbb{Z} -to- \mathbb{R} adaptation for all modalities $\mathbb{U}^{\downarrow}, \mathbb{S}^{\downarrow}, \mathbb{R}, \mathbb{T}$ such that any discrete-endpoint $\phi \in \text{bMTL}$ using only modalities in $\{\mathbb{U}^{\downarrow}, \mathbb{S}^{\downarrow}, \mathbb{R}, \mathbb{T}\}$ is c.u.i.s. with respect to it and $\sigma_{\delta,z}$. Then, $c \models_{\mathbb{R}} \eta_{\delta}^{\mathbb{Z}}[\phi]$ implies $c \models_{\mathbb{R}} v^{\mathbb{Z}}[\phi]$ for any behavior $c \in \mathcal{BPR}_{\delta}$.*

3.4 Generalizations

Theorem 15 proved that bMTL is sampling invariant. We claimed previously that bMTL is an MTL fragment of significant expressiveness; the specification examples in [Furia et al. 2008a; 2008b] demonstrate this in practice. Nevertheless, we are still interested in investigating to what extent Theorem 15 can be generalized to larger classes of MTL formulas. More precisely, given that Theorems 17 and 18 showed that canonical adaptations are optimal, we look for larger MTL fragments that are still sampling invariant with respect to $\eta_{\delta}^{\mathbb{R}}$ and $\eta_{\delta}^{\mathbb{Z}}$. Thus, henceforth sampling invariance will always implicitly refer to sampling invariance with respect to $\eta_{\delta}^{\mathbb{R}}$ and $\eta_{\delta}^{\mathbb{Z}}$.

Let us start by illustrating the rather apparent fact that, for any sampling period δ , there exist MTL formulas that are not s.i. with respect to δ .

EXAMPLE 19 A FORMULA NOT C.U.S.. *For an arbitrary sampling period δ , let us consider formula $\psi_{\delta} = \text{Som}(\square_{<\delta}(\mathbf{p}))$ and show that it is not c.u.s. with respect to δ . Consider any $c \in \mathcal{BPR}_{\delta}$ such that $\mathbf{p} \in c(t)$ iff $t \in V$ for some interval V such that $\delta < |V| < 2\delta$; clearly, $c \models_{\mathbb{R}} \psi_{\delta}$. However, for any z such that $\delta + z + \delta \lceil (\inf V - z)/\delta \rceil > \sup V$, \mathbf{p} holds at one unique sampling instant over $\sigma_{\delta,z}[c]$ (see Figure 3). Hence, $\sigma_{\delta,z}[c] \not\models_{\mathbb{Z}} \eta_{\delta}^{\mathbb{R}}[\psi_{\delta}]$, where $\eta_{\delta}^{\mathbb{R}}[\psi_{\delta}]$ corresponds to $\text{Som}(\square_{<1}(\mathbf{p}))$, because $\square_{<1}(\mathbf{p})$ requires \mathbf{p} to hold over two adjacent time instants. From the fact that the choice of origin z is arbitrary in the definition of sampling invariance (Definition 14), it follows that ψ_{δ} is not c.u.s.*




Fig. 3. Behavior c and its sampling $\sigma_{\delta,z}[c]$.

3.4.1 *Shiftable Formulas.* Examples 19 and 10 suggest a straightforward criterion to identify non-flat MTL formulas that are c.u.s.: if non-Berkeleyness can be “lifted” from propositional letters to the truth value of some nested sub-formula λ , then the nesting formula containing λ as a sub-formula can be flattened to one that is equi-satisfiable over non-Berkeley behaviors and does not introduce additional constraints. To formalize this notion, we introduce the following.¹⁰

DEFINITION 20 ϵ -SHIFTABILITY. *Formula ϕ is ϵ -shiftable, for some positive real ϵ , iff $b_\phi \in \mathcal{BPR}_\epsilon$ holds for all $b \in \mathcal{BPR}_\epsilon$. If ϕ is ϵ -shiftable for any ϵ , it is called shiftable.*

Shiftableity provides a straightforward condition to determine larger MTL subsets that are c.u.s. and c.u.i.s., as the following theorem shows.¹¹

THEOREM 21. *Let ψ be a shiftable formula.*

- (1) $\{\psi\}$ -bMTL and bMTL are equi-satisfiable over \mathcal{BPR}_δ for any δ .
- (2) If ψ and $\neg\psi$ are c.u.s., and $\neg\eta_\delta^R[\psi] \equiv \eta_\delta^R[\neg\psi]$ for all δ , then all $\{\psi\}$ -bMTL formulas are c.u.s.
- (3) If, for all $\psi' \in \eta_\delta^{Z^{-1}}[\psi]$ (where $\eta_\delta^{Z^{-1}}$ is the preimage of η_δ^Z), ψ' and $\neg\psi'$ are c.u.i.s., and $\neg\eta_\delta^Z[\psi'] \equiv \eta_\delta^Z[\neg\psi']$ for all δ , then all $\eta_\delta^{Z^{-1}}[\psi]$ -bMTL formulas are c.u.i.s.

PROOF. (1). Every $\{\psi\}$ -bMTL formula ϕ can be flattened into a bMTL formula $\bar{\phi}$ by introducing an auxiliary propositional letter \mathbf{a} that replaces every occurrence of ψ and is declared to be logically equivalent to ψ itself. Since ψ is shiftable, $b \in \mathcal{BPR}_\delta$ implies $b_\psi^{\psi \setminus \mathbf{a}} \in \mathcal{BPR}_\delta$; also, $b_\psi^{\psi \setminus \mathbf{a}} \in \mathcal{BPR}_\delta$ implies $b \in \mathcal{BPR}_\delta$ because b has no more transition points than $b_\psi^{\psi \setminus \mathbf{a}}$. Hence ϕ and $\bar{\phi}$ are equi-satisfiable.

(2). Let ϕ be any formula in $\{\psi\}$ -bMTL, and consider a behavior $b \in \mathcal{BPT}_\delta$ such that $b \models_{\mathbb{R}} \phi$. Let $\bar{\phi}$ denote the bMTL formula obtained by replacing every occurrence of ψ in ϕ by a fresh proposition $\mathbf{a} \in \mathcal{P}$, and let ϕ° be $\bar{\phi} \wedge (\mathbf{a} \Leftrightarrow \psi)$. Clearly, $b_\psi^{\psi \setminus \mathbf{a}} \models_{\mathbb{R}} \phi^\circ$, and $b_\psi^{\psi \setminus \mathbf{a}} \in \mathcal{BPT}_\delta$ as we showed in (1). Since $\bar{\phi}$ is flat, it is c.u.s. from Theorem 15; hence $b' \models_{\mathbb{Z}} \eta_\delta^R[\bar{\phi}]$ where $b' = \sigma_{\delta,z}[b_\psi^{\psi \setminus \mathbf{a}}]$. Notice that: $\eta_\delta^R[\mathbf{a} \Leftrightarrow \psi] = \eta_\delta^R[\mathbf{a} \wedge \psi \vee \neg\mathbf{a} \wedge \neg\psi]$ can be written as $\mathbf{a} \wedge \eta_\delta^R[\psi] \vee \neg\mathbf{a} \wedge \neg\eta_\delta^R[\psi]$. From the c.u.s. of both ψ and $\neg\psi$ and the fact that $\neg\eta_\delta^R[\psi] \equiv \eta_\delta^R[\neg\psi]$ we have $\eta_\delta^R[\mathbf{a} \Leftrightarrow \psi] = \mathbf{a} \Leftrightarrow \eta_\delta^R[\psi]$ and $b' \models_{\mathbb{Z}} \mathbf{a} \Leftrightarrow \eta_\delta^R[\psi]$. Let ϕ' be obtained from $\eta_\delta^R[\bar{\phi}]$ by substituting every occurrence of \mathbf{a} with $\eta_\delta^R[\psi]$. Hence, $b' \models_{\mathbb{Z}} \phi'$, which proves that ϕ is c.u.s.

(3). All similar to (2), by noticing that $\eta_\delta^Z[\psi'] = \psi$ for all $\psi' \in \eta_\delta^{Z^{-1}}[\psi]$ by definition of preimage. \square

3.4.2 *LTL is Nestable.* Theorem 21 is applicable to a significant class of MTL formulas, namely qualitative formulas. Indeed, LTL formulas are shiftable.¹²

LEMMA 22. *All LTL formulas are shiftable.*

¹⁰This notion is very similar to the notion of stability introduced in [Rabinovich 2003].

¹¹Recall the definition of Υ -bMTL at the end of Section 2.2.1.

¹²Note that non-strictness of LTL operators is necessary to have shiftableity.

PROOF. Let us consider any non-Berkeley behavior $b \in \mathcal{BPT}_\delta$ and any LTL formula ϕ . By induction, we prove that $\tau(b_\phi) \subseteq \tau(b)$ which subsumes the lemma.

The base case $\phi = \mathbf{p}$ is trivial. The case $\phi = \neg\phi_1$ follows from the inductive hypothesis $\tau(b_{\phi_1}) \subseteq \tau(b)$ because $\tau(b_{\phi_1}) = \tau(b_{\neg\phi_1})$.

Let us consider $\phi = \mathbf{U}(\phi_1, \phi_2)$; we consider $b' = b_\phi|_\phi$ and prove that $\tau(b') \subseteq \tau(b)$. To this end, let us first take any t such that $b'(t) \models_{\mathbf{R}} \phi$; hence $b(d) \models_{\mathbf{R}} \phi_2$ for some $d \geq t$, and $b(u) \models_{\mathbf{R}} \phi_1$ for all $u \in [t, d)$. The semantics of the qualitative *until* entails that $b'(t') \models_{\mathbf{R}} \phi$ holds for all $t \leq t' \leq d$. Then, ϕ cannot become false after d until ϕ_2 or ϕ_1 becomes false; similarly, ϕ cannot become false before t unless ϕ_1 becomes false. A dual argument shows that the same holds for $\neg\phi$. This establishes that $\tau(b') \subseteq \tau(b)$.

The last case that has to be considered is $\phi = \phi_1 \wedge \phi_2$. This is straightforward from the inductive hypothesis on ϕ_1 and ϕ_2 : $\tau(b_{\phi_1}) \subseteq \tau(b)$ and $\tau(b_{\phi_2}) \subseteq \tau(b)$. In addition, $\tau(b_{\phi_1 \wedge \phi_2}) \subseteq \tau(b_{\phi_1}) \cup \tau(b_{\phi_2})$ from the semantics of conjunction, hence $\tau(b_\phi) \subseteq \tau(b)$. It is simple to check that $b_\phi \in \mathcal{BPT}_\delta$ as well, because no left- and right- discontinuity can occur in b_ϕ as a result of applying conjunction. \square

Based on the previous lemma, the following corollary of Theorem 21 shows that any LTL qualitative formula can be nested within \mathbf{bMTL} formulas without losing c.u.s.

COROLLARY 23. *All LTL- \mathbf{bMTL} formulas are c.u.s.*

PROOF. The proof goes by induction on the nesting depth (i.e., the maximum number of nested modalities) of LTL formulas. For any integer $k > 0$, let LTL^k denote the set of all LTL formulas of nesting depth k .

The base case is for any flat LTL formula $\psi_1 \in \text{LTL}^1$. ψ_1 is shiftable from Lemma 22; ψ_1 and $\neg\psi_1$ are both c.u.s. from Theorem 15 (because $\neg\psi_1$ can also be written as a flat formula); one can check that $\neg\eta_\delta^{\mathbf{R}}[\psi_1] \equiv \eta_\delta^{\mathbf{R}}[\neg\psi_1]$ by pushing negations down to propositional letters. So all LTL^1 - \mathbf{bMTL} formulas are c.u.s. from Theorem 21.

Let now $\psi_k \in \text{LTL}^k$ be any LTL formula of nesting depth $k > 1$. ψ_k is shiftable from Lemma 22; ψ_k and $\neg\psi_k$ are both c.u.s., because they can both be written as LTL^{k-1} - \mathbf{bMTL} formulas, all of which are c.u.s. by inductive hypothesis; one can also check that $\neg\eta_\delta^{\mathbf{R}}[\psi_k] \equiv \eta_\delta^{\mathbf{R}}[\neg\psi_k]$ by pushing negations down to propositional letters and using the inductive hypothesis again. So all LTL^k - \mathbf{bMTL} formulas are c.u.s. from Theorem 21. \square

A similar corollary for c.u.i.s. cannot be obtained along the same lines, due to the transformation of *until* and its dual *release* under the canonical adaptation $\eta_\delta^{\mathbf{Z}}$.

4. VERIFICATION VIA SAMPLING

The notion of sampling invariance defines rigorously the connection between the non-Berkeley dense-time semantics and the discrete-time semantics of MTL, under the sampling relationship. On the one hand, this allows the formal description — by means of temporal logic formulas — of systems where dense-time and discrete-time components evolve in parallel, and communicate through a sampler. In addition, the theory of the previous sections can spawn several derived results that facilitate the analysis of real-time systems at the interface between discrete and dense time.

For instance, the notion of sampling can be used to describe system refinements from a “physical” dense-time model — close to a “real-world” physical description — to a more abstract discrete-time model — which is implementable on digital hardware.

This section investigates another significant application of the notion of sampling and sampling invariance. Namely, it builds a verification technique for dense-time MTL based on discretization. The intuition is that, in order to analyze the behaviors induced by a set of dense-endpoint \mathfrak{b} MTL formulas, their discrete-time samplings are analyzed instead. The results about sampling invariance allow us to move the results of the discrete-time analysis back to the dense-time domain, under some restrictions.

The following Section 4.1 shows how to build discrete-time under- and over-approximations of any \mathfrak{b} MTL formula. The over-approximation embodies discrete-time behaviors that are preserved into dense time, whereas the under-approximation represents discrete-time counter-examples that are preserved into dense time. Together, they allow a partial reduction of dense-time satisfiability for \mathfrak{b} MTL over non-Berkeley behaviors to dense-time MTL satisfiability. In order to perform system verification — i.e., checking if a given system satisfies certain putative properties — the under- and over-approximations of formulas can be combined to build two instances of the verification problem in the form of two validity checking problems for discrete-endpoint MTL formulas. This procedure is shown in Section 4.2. Finally, Section 4.3 comments on a few key issues of this verification procedure, in particular its strengths and weaknesses from a mostly practical viewpoint.

4.1 Under- and Over- Approximations

The over- and under-approximation functions $\Omega_\delta, \mathcal{O}_\delta$ are mappings from dense-endpoint \mathfrak{b} MTL formulas to discrete-endpoint \mathfrak{b} MTL formulas, parametric with respect to a sampling period δ . Given a \mathfrak{b} MTL formula ϕ , $\Omega_\delta[\phi]$ and $\mathcal{O}_\delta[\phi]$ retain some properties of the discrete-time *samplings* of the dense-time behaviors in \mathcal{BPR}_δ satisfying ϕ . Correspondingly, it is possible to infer the validity of ϕ over dense time from the validity of its approximations. For reasons that will become apparent shortly, $\Omega_\delta[\phi]$ is named *under-approximation* of ϕ and $\mathcal{O}_\delta[\phi]$ *over-approximation*. Unsurprisingly, $\Omega_\delta, \mathcal{O}_\delta$ are closely related to canonical adaptations $\eta_\delta^{\mathbb{R}}, \eta_\delta^{\mathbb{Z}}$; in particular the over-approximation is a sort of inverse of the mapping $\eta_\delta^{\mathbb{Z}}$. Their precise definition requires the introduction of the notion of granularity.

4.1.1 Granularity. For an MTL formula ϕ , let $\mathcal{I}_\phi = \{r_i/R_i\}_i$ be the set of all non-null, finite interval end-points appearing in ϕ and put in their irreducible form.¹³ The *granularity* ρ_ϕ of ϕ is defined as the pair: $\rho_\phi = (r_\phi, R_\phi) = (\gcd_i r_i, \text{lcm}_i R_i)$. Correspondingly, let us consider the set \mathcal{D}_ϕ of rationals:¹⁴

$$\mathcal{D}_\phi = \left\{ \frac{d}{D} \mid d \mid r_\phi \text{ and } R_\phi \mid D \right\}$$

It can be shown that, for any positive rational δ and $q \in \mathcal{I}_\phi$, q/δ is an integer iff $\delta \in \mathcal{D}_\phi$; i.e., \mathcal{D}_ϕ is the set of sampling periods δ such that any interval bound

¹³Recall that all finite endpoints are rationals (Section 2.2.1).

¹⁴Recall that $a \mid b$ denotes that b is an integer multiple of a .

in ϕ is an integer when divided by δ . Notice that \mathcal{D}_ϕ has a maximum (given by r_ϕ/R_ϕ) but no minimum. Finally, for a set of formulas Φ , \mathcal{D}_Φ is defined as $\mathcal{D}_{\widehat{\phi}}$ where $\widehat{\phi} \triangleq \bigwedge_{\varphi \in \Phi} \varphi$.

4.1.2 Under-Approximation. The under-approximation function Ω_δ maps dense-endpoint MTL formulas to discrete-endpoint MTL formulas such that the non-validity of the latter implies the non-validity of the former, over behaviors in \mathcal{BPT}_δ . More precisely, $\Omega_\delta[\phi]$ is defined only for MTL formulas such that δ is in \mathcal{D}_ϕ , where it coincides with $\eta_\delta^{\mathbb{R}}[\phi]$.

$$\begin{aligned} \Omega_\delta[\pi] &\triangleq \pi \\ \Omega_\delta\left[\mathbf{U}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{U}_{\lfloor l/\delta, u/\delta \rfloor}(\Omega_\delta[\phi_1], \Omega_\delta[\phi_2]) \\ \Omega_\delta\left[\mathbf{S}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{S}_{\lfloor l/\delta, u/\delta \rfloor}(\Omega_\delta[\phi_1], \Omega_\delta[\phi_2]) \\ \Omega_\delta\left[\mathbf{R}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{R}_{\lfloor l/\delta, u/\delta \rfloor}(\Omega_\delta[\phi_1], \Omega_\delta[\phi_2]) \\ \Omega_\delta\left[\mathbf{T}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{T}_{\lfloor l/\delta, u/\delta \rfloor}(\Omega_\delta[\phi_1], \Omega_\delta[\phi_2]) \\ \Omega_\delta[\phi_1 \wedge \phi_2] &\triangleq \Omega_\delta[\phi_1] \wedge \Omega_\delta[\phi_2] \\ \Omega_\delta[\phi_1 \vee \phi_2] &\triangleq \Omega_\delta[\phi_1] \vee \Omega_\delta[\phi_2] \end{aligned}$$

The following lemma justifies the name *under-approximation*.

LEMMA 24 UNDER-APPROXIMATION. *For any dense-endpoint bMTL formula ϕ , $\delta \in \mathcal{D}_\phi$, and $b \in \mathcal{BPZ}$: if $b \not\models_{\mathbb{Z}} \Omega_\delta[\phi]$ then for all $b' \in \mathcal{BPR}_\delta$ such that $\sigma_{\delta, z}[b'] = b$ it is $b' \not\models_{\mathbb{R}} \phi$.*

PROOF. ϕ is a dense-endpoint bMTL formula, hence it is c.u.s. from Theorem 15: for any $b \in \mathcal{BPR}_\delta$, if $b \models_{\mathbb{R}} \phi$ then $\sigma_{\delta, z}[b] \models_{\mathbb{Z}} \eta_\delta^{\mathbb{R}}[\phi]$. By taking the contrapositive, and by noticing that $\eta_\delta^{\mathbb{R}}$ and Ω_δ coincide when they are both defined, we have that for any $b \in \mathcal{BPR}_\delta$, if $\sigma_{\delta, z}[b] \not\models_{\mathbb{Z}} \Omega_\delta[\phi]$ then $b \not\models_{\mathbb{R}} \phi$. \square

4.1.3 Over-Approximation. The over-approximation function \mathbf{O}_δ maps dense-endpoint MTL formulas to discrete-endpoint MTL formulas such that the validity of the latter implies the validity of the former, over behaviors in \mathcal{BPT}_δ . More precisely, $\mathbf{O}_\delta[\phi]$ is defined only for MTL formulas such that δ is in \mathcal{D}_ϕ , where it is a pseudo-inverse of $\eta_\delta^{\mathbb{Z}}$.

$$\begin{aligned} \mathbf{O}_\delta[\pi] &\triangleq \pi \\ \mathbf{O}_\delta\left[\mathbf{U}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{U}_{\lfloor l/\delta+1, u/\delta-1 \rfloor}^\downarrow(\mathbf{O}_\delta[\phi_1], \mathbf{O}_\delta[\phi_2]) \\ \mathbf{O}_\delta\left[\mathbf{S}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{S}_{\lfloor l/\delta+1, u/\delta-1 \rfloor}^\downarrow(\mathbf{O}_\delta[\phi_1], \mathbf{O}_\delta[\phi_2]) \\ \mathbf{O}_\delta\left[\mathbf{R}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{R}_{\lfloor l/\delta-1, u/\delta+1 \rfloor}(\mathbf{O}_\delta[\phi_1], \mathbf{O}_\delta[\phi_2]) \\ \mathbf{O}_\delta\left[\mathbf{T}_{\langle l, u \rangle}(\phi_1, \phi_2)\right] &\triangleq \mathbf{T}_{\lfloor l/\delta-1, u/\delta+1 \rfloor}(\mathbf{O}_\delta[\phi_1], \mathbf{O}_\delta[\phi_2]) \\ \mathbf{O}_\delta[\phi_1 \wedge \phi_2] &\triangleq \mathbf{O}_\delta[\phi_1] \wedge \mathbf{O}_\delta[\phi_2] \\ \mathbf{O}_\delta[\phi_1 \vee \phi_2] &\triangleq \mathbf{O}_\delta[\phi_1] \vee \mathbf{O}_\delta[\phi_2] \end{aligned}$$

The following lemma justifies the name *over-approximation*.

LEMMA 25 OVER-APPROXIMATION. *For any dense-endpoint bMTL formula ϕ , $\delta \in \mathcal{D}_\phi$, and $b \in \mathcal{BPZ}$: if $b \models_{\mathbb{Z}} \text{O}_\delta[\phi]$ then for all $b' \in \mathcal{BPR}_\delta$ such that $\sigma_{\delta,z}[b'] = b$ it is $b' \models_{\mathbb{R}} \phi$.*

PROOF. If ϕ is a dense-endpoint bMTL formula, then $\text{O}_\delta[\phi]$ is a discrete-endpoint bMTL formula. Hence the latter is c.u.i.s. from Theorem 15: for any $b \in \mathcal{BPZ}$, if $b \models_{\mathbb{Z}} \text{O}_\delta[\phi]$ then $b' \models_{\mathbb{R}} \eta_\delta^{\mathbb{Z}}[\text{O}_\delta[\phi]]$ holds for all $b' \in \mathcal{BPR}_\delta$ such that $b = \sigma_{\delta,z}[b']$.

One can check that the dense-time validity of the formula $\eta_\delta^{\mathbb{Z}}[\text{O}_\delta[\phi]] \Rightarrow \phi$ is guaranteed by the definitions of $\eta_\delta^{\mathbb{Z}}$ and O_δ . In particular, $\eta_\delta^{\mathbb{Z}} \circ \text{O}_\delta$ is an identity for *release* (and *trigger*) operators with closed intervals. On the other hand, $\eta_\delta^{\mathbb{Z}} \circ \text{O}_\delta$ yields stronger formulas for *release* (and *trigger*) operators with open intervals and for *until* (and *since*) operators. The latter holds also from the fact that $\eta_\delta^{\mathbb{Z}} \left[\text{U}_{[l,u]}^\downarrow(\pi_1, \pi_2) \right]$ is $\text{U}_{((l-1)\delta, (u+1)\delta)}^\downarrow(\pi_1, \pi_2)$. It is easy to check that these properties of basic operators can be lifted to whole formulas by application of straightforward propositional identities on the negation normal form in which MTL formulas are expressed. In all, $b' \models_{\mathbb{R}} \eta_\delta^{\mathbb{Z}}[\text{O}_\delta[\phi]]$ implies $b' \models_{\mathbb{R}} \phi$. \square

4.2 MTL Verification

In the formal timed setting, verification consists in checking whether all behaviors generated by a system model (usually called *specification*) satisfy some given putative property (usually called *requirements*) [Heitmeyer and Mandrioli 1996]. Assume that both the specification and the requirements are formalized as MTL formulas ϕ_{sys} and ϕ_{prop} , respectively. Verification of ϕ_{sys} against ϕ_{prop} is equivalent to checking the validity of the dense-endpoint MTL formula $\phi_{\text{verif}} = \text{Alw}(\phi_{\text{sys}}) \Rightarrow \text{Alw}(\phi_{\text{prop}})$. If ϕ_{verif} is valid, any behavior of the system also respects the requirements; i.e., we have checked that $\llbracket \phi_{\text{sys}} \rrbracket_{\mathbb{R}} \subseteq \llbracket \phi_{\text{prop}} \rrbracket_{\mathbb{R}}$. On the contrary, if ϕ_{verif} is not valid, there exists at least one behavior of the system that violates the requirements; i.e., $\llbracket \phi_{\text{sys}} \rrbracket_{\mathbb{R}} \cap \llbracket \neg \phi_{\text{prop}} \rrbracket_{\mathbb{R}}$ is not empty so $\llbracket \phi_{\text{sys}} \rrbracket_{\mathbb{R}} \not\subseteq \llbracket \phi_{\text{prop}} \rrbracket_{\mathbb{R}}$.

In this section, we describe a verification algorithm that is applicable to specifications and requirements in bMTL over non-Berkeley dense-time behaviors.

The algorithm is based on the following.

PROPOSITION 26 MODEL APPROXIMATIONS. *For any bMTL formulas ϕ_1, ϕ_2 , and for any $\delta \in \mathcal{D}_{\{\phi_1, \phi_2\}}$:*

- (1) *if $\text{Alw}(\Omega_\delta[\phi_1]) \Rightarrow \text{Alw}(\text{O}_\delta[\phi_2])$ is \mathbb{Z} -valid, then $\text{Alw}(\phi_1) \Rightarrow \text{Alw}(\phi_2)$ is \mathbb{R}^δ -valid;*
- (2) *if $\text{Alw}(\text{O}_\delta[\phi_1]) \Rightarrow \text{Alw}(\Omega_\delta[\phi_2])$ is not \mathbb{Z} -valid, then $\text{Alw}(\phi_1) \Rightarrow \text{Alw}(\phi_2)$ is not \mathbb{R}^δ -valid.*

PROOF. (1). Let $\delta \in \mathcal{D}_{\{\phi_1, \phi_2\}}$. Assume that $\text{Alw}(\Omega_\delta[\phi_1]) \Rightarrow \text{Alw}(\text{O}_\delta[\phi_2])$ is \mathbb{Z} -valid. That is, for all $b \in \mathcal{BPZ}$ it is $b \not\models_{\mathbb{Z}} \Omega_\delta[\phi_1]$ or $b \models_{\mathbb{Z}} \text{O}_\delta[\phi_2]$. From Lemmas 25 and 24, this implies that for all $b \in \mathcal{BPZ}$, for all $b' \in \mathcal{BPR}_\delta$ such that $\sigma_{\delta,z}[b'] = b$, it is either $b' \not\models_{\mathbb{R}} \phi_1$ or $b' \models_{\mathbb{R}} \phi_2$. Since $\sigma_{\delta,z}$ is total, for any $b' \in \mathcal{BPR}_\delta$ there exists a $b \in \mathcal{BPZ}$ such that $\sigma_{\delta,z}[b'] = b$. We conclude that for all $b' \in \mathcal{BPR}_\delta$, either $b' \not\models_{\mathbb{R}} \phi_1$ or $b' \models_{\mathbb{R}} \phi_2$; i.e., $\text{Alw}(\phi_1) \Rightarrow \text{Alw}(\phi_2)$ is \mathbb{R}^δ -valid.

Proof of (2) is obtainable from the proof of (1) by duality. \square

4.2.1 *Verification Algorithm.* Proposition 26 suggests to introduce the following notation. Given a set of formulas $\Phi_{\text{sys}} = \{\phi_{\text{sys}}^i\}_i$ such that $\phi_{\text{sys}} = \bigwedge_i \text{Alw}(\phi_{\text{sys}}^i)$

represents a formal model of the system, and a formula ϕ_{prop} that represents a formal statement of the requirements, let us define the discrete-endpoint formulas:

$$\begin{aligned}\phi^{\text{O}} &\triangleq \bigwedge_i \text{Alw}(\Omega_\delta[\phi_{\text{sys}}^i]) \Rightarrow \text{Alw}(\text{O}_\delta[\phi_{\text{prop}}]) \\ \phi^{\Omega} &\triangleq \bigwedge_i \text{Alw}(\text{O}_\delta[\phi_{\text{sys}}^i]) \Rightarrow \text{Alw}(\Omega_\delta[\phi_{\text{prop}}])\end{aligned}$$

Let us call ϕ^{O} and ϕ^{Ω} *over-model* and *under-model* of the system, respectively, (in analogy with Lemmas 25 and 24) because the former preserves validity and the latter non-validity.

A verification algorithm for systems and properties specified as dense-endpoint bMTL formulas can be formalized as follows, where Z-VALID? is a validity-checking procedure for discrete-endpoint MTL formulas.

bMTL-VERIFY($\delta : \mathbb{R}_{>0}$, $\Phi_{\text{sys}} = \{\phi_{\text{sys}}^i\}_i$, $\phi_{\text{prop}} : \text{bMTL}$) : $\{\top, \perp, \text{FAIL}\}$

```

1  assume  $\delta \in \mathcal{D}_{\Phi_{\text{sys}} \cup \{\phi_{\text{prop}}\}}$ 
2   $\phi^{\text{O}} \leftarrow \bigwedge_i \text{Alw}(\Omega_\delta[\phi_{\text{sys}}^i]) \Rightarrow \text{Alw}(\text{O}_\delta[\phi_{\text{prop}}])$ 
3   $\phi^{\Omega} \leftarrow \bigwedge_i \text{Alw}(\text{O}_\delta[\phi_{\text{sys}}^i]) \Rightarrow \text{Alw}(\Omega_\delta[\phi_{\text{prop}}])$ 
4  if Z-VALID? $(\phi^{\text{O}})$  ▷  $\phi^{\text{O}}$  valid over discrete time?
5     then return  $\top$  ▷ verification over  $\mathcal{BPT}_\delta$  successful
6     else if  $\neg$ Z-VALID? $(\phi^{\Omega})$  ▷  $\phi^{\Omega}$  not valid over discrete time?
7         then return  $\perp$  ▷ verification over  $\mathcal{BPT}_\delta$  not successful
8         else return FAIL ▷ cannot conclude any verification result

```

The correctness of the algorithm follows directly from Proposition 26, keeping in mind that $b \models_{\top} \text{Alw}(\psi_1) \wedge \text{Alw}(\psi_2)$ iff $b \models_{\top} \text{Alw}(\psi_1)$ and $b \models_{\top} \text{Alw}(\psi_2)$.

4.2.2 Incompleteness. A verification algorithm is *complete* if, for any input, it terminates with a conclusive result about whether the given requirements ϕ_{prop} are indeed a property of the system ϕ_{sys} or not.

The verification algorithm for bMTL we provided above is *incomplete*, as it can fail to provide a conclusive answer about whether ϕ_{prop} is indeed a property of all behaviors of the system ϕ_{sys} . The incompleteness is two-fold. First, the algorithm does not consider *all* dense-time behaviors \mathcal{BPR} , but only those in \mathcal{BPT}_δ , i.e., “slow” with respect to some chosen sampling period δ . Hence, it may be that ϕ_{prop} does not hold for some “real” behavior of the system which is “fast”, i.e., for some behavior in $[[\phi_{\text{sys}}]]_{\mathbb{R}} \setminus [[\phi_{\text{sys}}]]_{\mathbb{R}}^\delta$. Second, the under- and over-model $\phi^{\Omega}, \phi^{\text{O}}$ are in general non-equivalent discrete-endpoint formulas. Hence, it is possible that ϕ^{O} is not valid and ϕ^{Ω} is valid; if this is the case no conclusion about the verification of the system can be drawn.

Since the algorithm is parametric with respect to δ , smaller values of δ can be tried in order to avoid the incompleteness hurdle. Changing the value of δ affects the verification problem in two ways: more (“faster”) behaviors are considered for verification, and new under- and over-models are generated that represent a “finer-grain” discretization of the original problem. These two aspects interact in subtle ways because they change the verification problem from two opposite sides. By combining them, one may expect to achieve at least the following partial notion of

completeness: if ϕ_{prop} is a property of ϕ_{sys} over behaviors in \mathcal{BPR}_δ for *some* choice of δ , then there exists a suitable choice of δ such that $\text{bMTL-VERIFY}(\delta, \phi_{\text{sys}}, \phi_{\text{prop}})$ returns \top ; and conversely when ϕ_{prop} is not a property of ϕ_{sys} . Unfortunately, the following example shows that even this weaker notion of completeness is not achieved by the algorithm.

EXAMPLE 27 INCOMPLETENESS OF THE ALGORITHM. *Consider a simple set of behaviors completely described by formulas in Table II. It should be clear that all behaviors $b \in \llbracket \text{Alw}(\phi_{\text{sys}}^1) \wedge \text{Alw}(\phi_{\text{sys}}^2) \rrbracket_{\mathbb{R}}$ of the system are such that \mathbf{p} holds on some interval $V = \langle t, +\infty \rangle$ and $\neg \mathbf{p}$ holds on the complement interval $\mathbb{R} \setminus V$ (which is unbounded to the left). Hence, any such b satisfies property ϕ_{prop} and is in \mathcal{BPR}_δ for any δ .*

ϕ_{sys}^1	\triangleq	$\text{Som}(\mathbf{p}) \wedge \text{Som}(\neg \mathbf{p})$
ϕ_{sys}^2	\triangleq	$\mathbf{p} \Rightarrow \Box(\mathbf{p})$
ϕ_{prop}	\triangleq	$\mathbf{p} \Rightarrow \Diamond_{=1}(\mathbf{p})$

Table II. Φ_{sys} and ϕ_{prop} .

Table III shows the over- and under-models of this system for any $\delta \in \mathcal{D}_{\Phi_{\text{sys}} \cup \{\phi_{\text{prop}}\}} = \{1/k \mid k \in \mathbb{N}_{>0}\}$, after some simplifications (in particular $\text{O}_\delta[\phi_{\text{sys}}^2] = \neg \mathbf{p} \vee \Box_{[-1, +\infty]}(\mathbf{p})$ is equivalent to the formula in Table III under the global satisfiability semantics). It is simple to check that, for any value of δ , the over-model ϕ^{O} is not valid because $\text{Alw}(\text{O}_\delta[\phi_{\text{prop}}])$ contradicts $\text{Alw}(\Omega_\delta[\phi_{\text{sys}}^1])$. Also for any value of δ the under-model ϕ^{Ω} is vacuously valid because $\text{Alw}(\text{O}_\delta[\phi_{\text{sys}}^1])$ is inconsistent with $\text{Alw}(\text{O}_\delta[\phi_{\text{sys}}^2])$. In all, we cannot verify our system with our algorithm, no matter what value of sampling period we choose.

$\Omega_\delta[\phi_{\text{sys}}^1] = \text{Som}(\mathbf{p}) \wedge \text{Som}(\neg \mathbf{p})$	$\text{O}_\delta[\phi_{\text{sys}}^1] = \text{Som}(\mathbf{p}) \wedge \text{Som}(\neg \mathbf{p})$
$\Omega_\delta[\phi_{\text{sys}}^2] = \mathbf{p} \Rightarrow \Box(\mathbf{p})$	$\text{O}_\delta[\phi_{\text{sys}}^2] = \text{Alw}(\mathbf{p}) \vee \text{Alw}(\neg \mathbf{p})$
$\Omega_\delta[\phi_{\text{prop}}] = \mathbf{p} \Rightarrow \Diamond_{=k}(\mathbf{p})$	$\text{O}_\delta[\phi_{\text{prop}}] = \neg \mathbf{p}$

Table III. Under- and over-models of $\Phi_{\text{sys}}, \phi_{\text{prop}}$ for $\delta = 1/k$.

In spite of its incompleteness, in the next section we discuss why the verification algorithm can still provide practically very useful results.

4.3 Discussion

In related work, we proved that MTL is fully decidable over dense-time non-Berkeley behaviors \mathcal{BPR}_δ for any δ [Furia and Rossi 2008], with the same worst-case complexity as discrete-time MTL; hence an *incomplete* decision procedure may seem impractical. In this section we demonstrate that this is not the case, and we discuss how the impact of incompleteness can be limited in practice with the application of a few good practices.

First of all, the decision procedure for MTL over \mathcal{BPR}_δ — the only one currently available [Furia and Rossi 2008] — relies on a rather exotic decision procedure,

which translates MTL to a family of uncommon decidable real-time temporal logics introduced by Hirshfeld and Rabinovich [Hirshfeld and Rabinovich 2004]. The decision procedures for such logics have never been implemented, and seem quite complex in practice. More generally, the practical high complexity of deciding temporal logics over dense-time domains is witnessed not only by theoretical results, but also by the current scarcity of state-of-the-art tools that implement such decision procedures. Even the well-known real-time temporal logic MITL, whose decidability over dense time is known since the seminal work of Alur, Feder, and Henzinger [Alur et al. 1996], still lacks an implementation, despite the recent efforts towards simplifying its decision procedure [Hirshfeld and Rabinovich 2005; Maler et al. 2006].

Compare this unsatisfactory picture to the vastly different scenario of (real-time) temporal logics over discrete time, where a significant number of off-the-shelf efficient verification tools are available (e.g., [Pradella et al. 2007; Bianculli et al. 2007; Pradella et al. 2003; Cimatti et al. 2002; De Wulf et al. 2009] just to mention a few for LTL/MTL). This suggests that a dense-time verification procedure based on discretization is very appealing from a practical viewpoint, because it can be implemented easily and it can rely on solid and scalable implementations. In fact, in related work [Furia et al. 2008a; 2008b; Bersani et al. 2009] we presented the straightforward implementation of the verification procedure described in this section, and we demonstrated its practical efficiency with a few non-trivial verification examples.

The same examples also show that the flat fragment of MTL retains (under the global satisfiability semantics) a significant expressive power, suitable to formalize typical behaviors of real-time systems. For example, it is possible to describe runs of arbitrary timed automata or bounded time Petri nets over non-Berkeley behaviors. The formalization in flat MTL of these complex abstract machines is far from straightforward and requires a careful analysis to avoid inconsistencies. However, the experience of [Furia et al. 2008a; 2008b; Bersani et al. 2009] can be leveraged and extended to similar systems described by means of the notions of state and transition.

Even the *incompleteness* of our verification algorithm turns out not to be too large a handicap in practice. More precisely, the fact that equivalent dense-endpoint formulas can yield nonequivalent discrete-time under- or over-approximations can be turned into an advantage: with some additional effort in writing the dense-time model of our system, we can often express it in a form whose over- and under-models are unaffected by incompleteness. This effort can in general be non-trivial, but it can give very good practical results nonetheless. The following example provides a few in-the-small demonstrations of our claims, whereas more complex cases have been introduced elsewhere [Furia et al. 2008b; Bersani et al. 2009].

EXAMPLE 28. *Let us go back to Example 27 and change formula ϕ_{sys}^2 into $\psi_{\text{sys}}^2 \triangleq \mathbf{p} \Rightarrow \square_{\geq \delta}(\mathbf{p})$, according to the chosen sampling period δ . A little reasoning should convince us that $\text{Alw}(\phi_{\text{sys}}^2)$ is equivalent to $\text{Alw}(\psi_{\text{sys}}^2)$ over behaviors in \mathcal{BPR}_δ : if \mathbf{p} holds at some time t as well as over the left-closed interval $t \oplus [\delta, +\infty)$, it cannot be false anywhere in $(t, t + \delta)$ because this would violate the hypothesis of non-Berkeleyness for the given δ . Let us take our system model to be $\Phi_{\text{sys}} = \{\phi_{\text{sys}}^1, \psi_{\text{sys}}^2\}$,*

and let us build its over-model $\Phi_{\text{sys}}^{\text{O}}$. Notice that $\text{O}_{\delta}[\psi_{\text{sys}}^2]$ can be computed as $\mathbf{p} \Rightarrow \Box(\mathbf{p})$; unlike $\text{O}_{\delta}[\phi_{\text{sys}}^2]$, this is an accurate discrete-time rendition of the dense-time model. It is now possible to prove that ϕ^{O} is \mathbb{Z} -valid for any $\delta = 1/k$, which verifies our system over dense time.

Let us now turn our attention to property ϕ_{prop} in Example 27. It should be apparent that its over-approximation $\text{O}_{\delta}[\phi_{\text{prop}}] = \neg \mathbf{p}$ is very unsatisfactory, and it is unlikely to yield valuable results when used in an under-model. Consider however formula $\phi'_{\text{prop}} \triangleq \mathbf{p} \Rightarrow \Box_{=1}(\mathbf{p})$; ϕ'_{prop} is trivially equivalent to ϕ_{prop} . However, its over-approximation is the much more reasonable $\mathbf{p} \Rightarrow \Box_{[k-1, k+1]}(\mathbf{p})$ which is non-trivially satisfiable for any $k > 1$.

5. RELATED WORK

The relationship between dense and discrete real-time semantics has been investigated by many authors. In this section we mention the approaches that are closest to ours, and we detail the most significant differences and relative merits.

The seminal paper by Henzinger, Manna, and Pnueli [Henzinger et al. 1992] is both the first and the best-known work dealing with the theme of dense vs. discrete real-time through the notion of *digitization*. Given the significance of this notion, Section 5.1 is devoted to a detailed summary of it, as well as to a comparison with sampling invariance. Section 5.2 succinctly describes other related work about the relation between dense and discrete time models for real-time formalisms. Finally, briefly widening the scope beyond real-time notations, the results of this paper seem to bear a connection with the classical theory of digital sampling (e.g., [Benedetto and Ferreira 2001]). Section 5.3 sketches a partly formal analysis of this alleged link.

5.1 Comparison with Digitization

Similarly to the notions of *sampling* and *sampling invariance* — introduced in Section 3 — the notions of *digitization* and *digitizability* [Henzinger et al. 1992] link dense- and discrete-time real-time semantics. The main purpose of digitization is to provide a means to reduce the verification problem from the richer dense-time semantics to the simpler discrete-time one. This section recalls the formal definition of digitization and digitizability and compares them against the notions of sampling, sampling invariance, and discrete-time approximations introduced in this paper.

There are two fundamental high-level differences between the frameworks of digitization and sampling; bridging them is necessary to carry out a formal comparison of the notions. First, our framework considers dense- and discrete-time *behaviors* as semantic structures, whereas digitization is defined for dense- and discrete-valued *timed words*. A timed word is a discrete sequence of timestamped events, such that every event is assumed to occur at the absolute time value of its timestamp. Second, sampling invariance is a syntactic notion (i.e., it is a property that applies to *formulas*), whereas digitizability is a semantic notion (i.e., it is a property that applies to *sets of timed words*). Let us introduce formally these ideas and the precise notions of digitization and digitizability.

DEFINITION 29 MTL TIMED WORD SEMANTICS. *An (infinite) timed word over*

\mathcal{P} is an ω -sequence $(\sigma_0, t_0)(\sigma_1, t_1) \cdots (\sigma_i, t_i) \cdots$ in $(\mathcal{P} \times \mathbb{T})^\omega$, such that the sequence of timestamps t_i is weakly monotonic and diverging. According to whether \mathbb{T} is a dense (typically $\mathbb{R}_{\geq 0}$) or discrete (typically \mathbb{N}) set, the timed words are named dense- or discrete-valued.

MTL semantics over timed words is defined as expected: given a timed word ρ , a position $i \in \mathbb{N}$, and an MTL formula ϕ , we write $\rho, i \models \phi$ iff ρ satisfies ϕ at position i . The definition of the modalities is: $\rho, i \models \mathbf{U}_I(\phi_1, \phi_2)$ iff there exists $j \geq i$ such that $t_j \in t_i \oplus I$, $\rho, j \models \phi_2$, and $\rho, k \models \phi_1$ for all $i \leq k < j$; and $\rho, i \models \mathbf{R}_I(\phi_1, \phi_2)$ iff for all $j \geq i$ such that $t_j \in t_i \oplus I$, it is $\rho, j \models \phi_2$ or $\rho, k \models \phi_1$ for some $i \leq k < j$. Then, $\rho \models \phi$ iff $\rho, i \models \phi$ for all $i \in \mathbb{N}$.¹⁵ Given a formula ϕ , $\langle\langle \phi \rangle\rangle_{\mathbb{T}}$ denotes the set $\{\rho \mid \rho \models \phi\}$ of \mathbb{T} -valued timed words that satisfy ϕ .

DEFINITION 30 DIGITIZATION AND DIGITIZABILITY. Given a timed word $\rho = \{(\sigma_i, t_i) \mid i \in \mathbb{N}\}$ and a fractional value $0 \leq \epsilon < 1$, the ϵ -digitization of ρ is defined as the discrete-valued timed word $[\rho]_\epsilon = \{(\sigma_i, [t_i]_\epsilon) \mid i \in \mathbb{N}\}$, where $[t]_\epsilon$ is $\lfloor t \rfloor$ if $t \leq \lfloor t \rfloor + \epsilon$, and $\lfloor t \rfloor$ otherwise. The digitization of a set of timed words Π is the set $[\Pi]$ of discrete-valued timed words defined as $\{[\rho]_\epsilon \mid \rho \in \Pi \text{ and } 0 \leq \epsilon < 1\}$, i.e., the set of all possible digitizations of words in Π .

A set of timed words Π is: (1) closed under digitization (c.u.d.) iff $\rho \in \Pi$ implies $[\rho] \subseteq \Pi$; (2) closed under inverse digitization (c.u.i.d.) iff $[\rho] \subseteq \Pi$ implies $\rho \in \Pi$; (3) digitizable iff it is c.u.d. and c.u.i.d. Correspondingly, an MTL formula ϕ is c.u.d., c.u.i.d., or digitizable, iff $\langle\langle \phi \rangle\rangle_{\mathbb{R}_{\geq 0}}$ is.

For digitizable properties, discrete-time verification completely captures dense-time verification; more precisely, if a system specification is closed under digitization, and the requirements are closed under inverse digitization, the problem of determining if the specification meets the requirements is perfectly reducible to the discrete-time case. However, it is difficult to characterize a significant syntactic subset of MTL formulas that are digitizable, and in fact only a few examples are given in [Henzinger et al. 1992]. Moreover, digitization exploits weakly-monotonic timed word to ensure that no dense-time event is lost when digitizing a dense-valued timed word; this is why no notion similar to non-Berkeleyness is introduced.

The following example shows that digitizability and sampling invariance define *incomparable classes* of MTL formulas, i.e., there exist sampling invariant non-digitizable formulas, as well as digitizable non sampling-invariant formulas. This demonstrates that the two notions have different angles, and it suggests that techniques for discrete-time verification of dense-time MTL formulas based on these two orthogonal notions may each have its own complementary strengths and weaknesses.

EXAMPLE 31. For $h \in \mathbb{N}_{>0}$, let Θ_h^{snd} be the bMTL formula $\mathbf{p} \Rightarrow \diamond_{<h}(\mathbf{q})$. Theorem 15 proves that Θ_h^{snd} is s.i. Let us show that Θ_h^{snd} is instead not c.u.d., hence neither digitizable. Take any timed word $\sigma = \cdots (\mathbf{p}, k)(\mathbf{q}, k+h-1+\mu)(\mathbf{q}, k+h+\mu) \cdots$ with $k \in \mathbb{N}$, $0 < \mu < 1$, and such that \mathbf{p} does not occur anywhere else. Any ϵ -

¹⁵The digitization paper [Henzinger et al. 1992] assumed an initial satisfiability semantics, but we adopt a global satisfiability semantics to allow a uniform comparison with sampling invariance (see Section 2.2.2); it should be clear that this is without loss of generality.

digitization of σ for $\epsilon < \mu$ has the form $[\sigma]_\epsilon = \dots(\mathbf{p}, k)(\mathbf{q}, k + h)(\mathbf{q}, k + h + 1)\dots$. Hence Θ_h^{snd} is not c.u.d. because $\sigma \models \Theta_h^{snd}$ but $[\sigma]_\epsilon \not\models \Theta_h^{snd}$ for any such ϵ .

For $h \in \mathbb{N}_{>0}$, let Θ_h^{dns} be the MTL formula $\text{Som}(\mathbf{p} \wedge \bigcirc(\neg\mathbf{p})) \wedge \psi_h$, where ψ_h has been defined in Example 19. It is not difficult to show that $\text{Som}(\mathbf{p} \wedge \bigcirc(\neg\mathbf{p}))$ is unsatisfiable in the timed word semantics, hence Θ_h^{dns} is trivially digitizable. Let us show that Θ_h^{dns} is instead not c.u.s., hence neither s.i. Take the same behavior $c \in \mathcal{BPR}_h$ of Example 19, where we further assume that V is a right-closed interval (see Figure 3). $c \models_{\mathbb{R}} \Theta_h^{dns}$ because Example 19 showed that $c \models_{\mathbb{R}} \psi_h$ and $\mathbf{p} \wedge \bigcirc(\neg\mathbf{p})$ holds at the right end-point of V . However, Example 19 also proved that $\sigma_{h,z}[c] \not\models_{\mathbb{Z}} \eta_h^{\mathbb{R}}[\psi_h]$, so $\sigma_{h,z}[c] \not\models_{\mathbb{Z}} \eta_h^{\mathbb{R}}[\Theta_h^{dns}]$ as well. Hence, Θ_h^{dns} is not c.u.s.

5.2 Other Work on the Relations between Dense and Discrete Time

The introduction of the notion of *digitization* has spawned much derivative work, where the notion is applied to various formalisms. Several authors considered digitization for automata-based real-time formalisms, especially timed automata [Bouajjani et al. 1994; Bošnački 1999; Maler and Pnueli 1995; Bozga et al. 1999; Beyer et al. 2003; Ouaknine and Worrell 2003; Clarke et al. 2007]. Others studied how the decidability and complexity of standard verification problems for *timed automata* (esp. reachability) change when moving from a dense- to a discrete-time semantics, such as in [Göllü et al. 1994; Krčál and Pelánek 2005]. Asarin, Maler, and Pnueli [Asarin et al. 1998] investigated instead to what extent *qualitative* behavior of digital circuits (which can in turn be modeled as timed automata [Maler and Pnueli 1995]) is preserved in a sampled discrete-time semantics. The focus of all these works is to determine to what extent the computationally simpler discrete-time semantics can be substituted for the dense-time semantics for automated verification.

The notion of digitization has been applied also to descriptive notations, such as real-time temporal logics and process algebras. In the latter category, Ouaknine studies digitization for *timed CSP* [Ouaknine 2002]; his main contribution is the proof that all CSP are closed under inverse digitization, hence they can be model-checked over dense time by considering just their discrete-time semantics.

Among temporal logics, the digitization of *duration calculus* (DC) and its variants has been studied in several works. Van Hung and Giang consider standard duration calculus and a slight generalization of digitization called *sampling* [Hung and Giang 1996]. Their work is focused on providing inference rules that allow one to infer the validity of dense-time formulas from the validity of sampled discrete-time formulas and *vice versa*. Another similarity with our approach is that they consider δ -stability: a constraint similar to non-Berkeleyness that relates the “speed” of signals and the sampling period δ . Unlike non-Berkeleyness δ -stability is asymmetric, in that whenever a proposition switches to true it must hold its truth value for more than δ time units, but it is not required to do so when it switches to false.

Pandya et al. also have applied the notion of digitization to DC, with the aim of developing efficient dense-time verification techniques based on discretization. Their overall approach consists of two parts, and it has been shown to be applicable to MTL as well [Pandya 2008]. In the first part [Chakravorty and Pandya 2003], the notion of digitization has been applied to IDL (Interval Duration Logic) a DC variant whose formulas are interpreted over timed words. Given that a syntactic characterization of closure under inverse digitization for IDL formulas is hard to

achieve, a new notion of *strong closure under inverse digitization* (SCID) is introduced. SCID eases the problem because it is straightforward to determine if an IDL formula is SCID, and SCID entails closure under inverse digitization in the standard sense. For formulas that are not SCID, approximations of formulas are introduced. In the second part [Pandya et al. 2007], the richer semantics of DC (based on behaviors) is reduced to the timed word semantics of IDL through two approximation mappings α^+ and α^- . α^+ and α^- play a role similar to our over- and under- approximations O_δ, Ω_δ , in that α^+ preserves non-validity and α^- preserves validity from the sampled to the dense-time semantics. Unsurprisingly, the resulting verification technique is incomplete, as DC is undecidable over dense time.

De Alfaro and Manna considered the problem of discretization for the predicate temporal logic TL [de Alfaro and Manna 1995]. Their results are based on the semantic notion of *finite variability*: informally, a formula ϕ is finitely variable if, for any timed word, one can find a refined “ground” timed word such that any subformula of ϕ has a constant truth value within any interval of the refined word. For finitely variable formulas over ground traces, the satisfaction relation of a formula ϕ in the dense-time semantics corresponds to that of $\Omega(\phi)$ in the discrete-time semantics (where Ω is a given translation function). Some sufficient syntactic conditions for a formula to achieve the finite variability requirement are introduced; based on these, a methodology for dense-time verification through refinement to discrete time is proposed.

Fainekos and Pappas [Fainekos and Pappas 2007; 2009] present a technique for testing specifications written in MITL (an MTL subset) against continuous-time signals by analyzing only discrete samplings of the signals. Their technique shares underlying motivations and ideas with ours, although the two approaches have complementary scopes: our results bridge the gap between the dense-time non-Berkeley semantics and the discrete-time semantics for MTL, whereas Fainekos and Pappas discover concrete and practical conditions under which the continuous-time behavior of a dynamical system can be analyzed by means of its discrete-time observations.

5.3 The Sampling Theorem

The sampling theorem [Benedetto and Ferreira 2001] states sufficient conditions for which no information loss occurs in the digital sampling of a continuous-time signal. A continuous-time *signal* s is a mapping $s : \mathbb{R} \rightarrow D$ where D is some — usually dense — codomain. B_s denotes the *bandwidth* of s , that is its highest frequency in s .¹⁶ Using the notation of Section 3, the *sampling* of s with sampling period δ is the discrete-time signal $\sigma_{\delta,0}[s]$. The sampling theorem states that s can be perfectly reconstructed from $\sigma_{\delta,0}[s]$ for any $\delta < 1/(2B_s)$.

A number of similarities between this fundamental theorem of signal theory and the results of this paper are apparent. In particular, the requirement on the relation between bandwidth and sampling period is reminiscent of the non-Berkeley requirement, so that the results of this paper might seem a consequence of the sampling theorem. Our dense-time behaviors \mathcal{BPR} can indeed be

¹⁶The highest frequency is defined as the largest nonzero value for which the Fourier transform $F[s]$ of s is non-zero.

modeled as continuous-time signals over range $[0, 2^{|\mathcal{P}|}]$. However all of them have *infinite bandwidth* because of the discontinuities corresponding to transition points, regardless of whether they are non-Berkeley or not. Hence the sampling theorem cannot strictly be applied to Boolean-valued signals. Nonetheless, a connection between the theory of sampling and the theory of this paper exists, as we demonstrate in the following.

EXAMPLE 32. Consider a simple unary alphabet $\{\mathbf{p}\}$ and a single behavior b such that \mathbf{p} holds over $\mathbb{R}_{>0}$ and does not hold over $\mathbb{R}_{<0}$ (we disregard the value of \mathbf{p} exactly at 0). b corresponds to the signal $s : \mathbb{R} \rightarrow [0, 1]$ defined as $s(t) = H(t)$ where H denotes the usual (Heaviside) unit step function (see Figure 4). b can also be




Fig. 4. Signals $s(t)$ (in gray) and $\hat{s}(t)$ (in black).

described perfectly by the MTL formula $\beta = \overleftarrow{\square}_{>0}(\neg\mathbf{p}) \wedge \square_{>0}(\mathbf{p})$ evaluated at the origin. The discrete-time MTL formula $\beta' = \eta_{\delta}^{\mathbb{R}}[\beta] = \overleftarrow{\square}_{\geq 1}(\neg\mathbf{p}) \wedge \square_{\geq 1}(\mathbf{p})$ characterizes discrete-time samplings of b according to our theory. β' can be seen as describing some dense-time behaviors in \mathcal{BPR}_{δ} through their samplings: all behaviors such that \mathbf{p} holds over $\mathbb{R}_{\geq \delta}$ and it does not hold over $\mathbb{R}_{\leq -\delta}$. Hence, the sampling has introduced an information loss in the formula about where exactly \mathbf{p} switches within $(-\delta, \delta)$. If we try to reconstruct s from its digital sampling according to the classical theory, we notice that we introduce a similar information loss. In fact, let $\hat{s} : \mathbb{R} \rightarrow [0, 1]$ be the continuous-time reconstruction of $\sigma_{\delta,0}[s]$ built with the Whittaker-Shannon interpolation formula, i.e., $\hat{s}(t) = \sum_{k \in \mathbb{Z}} \sigma_{\delta,0}[s](k) \text{sinc}((t - k\delta)/\delta)$. As it can be seen in Figure 4, \hat{s} coincides almost perfectly with s over $\mathbb{R}_{\leq -\delta} \cup \mathbb{R}_{\geq \delta}$ (the residual errors are only due to numerical approximations), whereas it deviates significantly within $(-\delta, \delta)$ due to the information loss introduced with sampling (it passes right through the origin only as a result of symmetry). In this sense information loss for Boolean-valued signals are similar in our theory for MTL and in classical sampling theory for signals.

6. CONCLUSION

In this paper, we presented an approach to relate dense-time MTL formulas to some discrete-time counterparts (and *vice versa*). We exploited the resulting relationship to define a technique for the verification through discretization of systems described as dense-time MTL formulas. The verification technique is inherently incomplete,

though in practice it has yielded promising results [Furia et al. 2008a; 2008b; Bersani et al. 2009].

In the future, we plan to apply the notion of sampling presented in this paper to the synthesis of software components of real-time systems from continuous-time specifications. We will also further investigate the properties of the verification technique presented in Section 4, in particular to better characterize, and possibly reduce, the scope of its incompleteness.

ACKNOWLEDGMENTS

We thank the anonymous reviewers for their detailed comments.

REFERENCES

- ALUR, R., FEDER, T., AND HENZINGER, T. A. 1996. The benefits of relaxing punctuality. *Journal of the ACM* 43, 1, 116–146.
- ALUR, R. AND HENZINGER, T. A. 1993. Real-time logics: Complexity and expressiveness. *Information and Computation* 104, 1, 35–77.
- ASARIN, E., MALER, O., AND PNUELI, A. 1998. On discretization of delays in timed automata and digital circuits. In *Proceedings of the 9th International Conference on Concurrency Theory (CONCUR'98)*, D. Sangiorgi and R. de Simone, Eds. Lecture Notes in Computer Science, vol. 1466. Springer-Verlag, 470–484.
- BENEDETTO, J. J. AND FERREIRA, P. J. S. G., Eds. 2001. *Modern Sampling Theory*. Birkäuser Boston.
- BERSANI, M. M., FURIA, C. A., PRADELLA, M., AND ROSSI, M. 2009. Integrated modeling and verification of real-time systems through multiple paradigms. In *Proceedings of the 7th IEEE International Conference on Software Engineering and Formal Methods (SEFM'09)*. IEEE Computer Society Press.
- BEYER, D., LEWERENTZ, C., AND NOACK, A. 2003. Rabbit: A tool for BDD-based verification of real-time systems. In *Proceedings of the 15th International Conference on Computer Aided Verification (CAV'03)*, W. A. H. Jr. and F. Somenzi, Eds. Lecture Notes in Computer Science, vol. 2725. Springer-Verlag, 122–125.
- BIANCULLI, D., MORZENTI, A., PRADELLA, M., SAN PIETRO, P., AND SPOLETINI, P. 2007. Trio2Promela: A model checker for temporal metric specifications. In *ICSE Companion*. 61–62.
- BOŠNAČKI, D. 1999. Digitization of timed automata. In *Proceedings of the 4th International Workshop on Formal Methods for Industrial Critical Systems (FMICS'99)*. 283–302.
- BOUAJJANI, A., ECHAHED, R., AND ROBBANA, R. 1994. Verifying invariance properties of timed systems with duration variables. In *Proceedings of the 3rd International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRFT'94)*. Lecture Notes in Computer Science, vol. 863. Springer-Verlag, 193–210.
- BOUYER, P., MARKEY, N., OUAKNINE, J., AND WORRELL, J. 2007. The cost of punctuality. In *Proceedings of the 22nd IEEE Symposium on Logic in Computer Science (LICS'07)*. IEEE Computer Society.
- BOZGA, M., MALER, O., AND TRIPAKIS, S. 1999. Efficient verification of timed automata using dense and discrete time semantics. In *Proceedings of the 10th Correct Hardware Design and Verification Methods Advanced Research Working Conference (CHARME'99)*, L. Pierre and T. Kropf, Eds. Lecture Notes in Computer Science, vol. 1703. Springer-Verlag, 125–141.
- CHAKRAVORTY, G. AND PANDYA, P. K. 2003. Digiziting interval duration logic. In *Proceedings of the 15th International Conference on Computer Aided Verification (CAV'03)*, W. A. Hunt, Jr. and F. Somenzi, Eds. Lecture Notes in Computer Science, vol. 2725. Springer-Verlag, 167–179.
- CIMATTI, A., CLARKE, E. M., GIUNCHIGLIA, E., GIUNCHIGLIA, F., PISTORE, M., ROVERI, M., SEBASTIANI, R., AND TACCHELLA, A. 2002. NuSMV 2: An open source tool for symbolic model checking. In *Proceeding of the 14th International Conference on Computer-Aided Verification (CAV'02)*. Lecture Notes in Computer Science, vol. 2404. Springer-Verlag, 359–364.

- CLARKE, E. M., LERDA, F., AND TALUPUR, M. 2007. An abstraction technique for real-time verification. In *Proceedings of the GM R&D Workshop on Next Generation Design and Verification Methodologies for Distributed Embedded Control System*.
- COMON, H. AND CORTIER, V. 2000. Flatness is not a weakness. In *Proceedings of the 14th Annual Conference of the EACSL on Computer Science Logic*. Lecture Notes in Computer Science, vol. 1862. Springer-Verlag, 262–276.
- DAMS, D. 1999. Flat fragments of CTL and CTL*: Separating the expressive and distinguishing powers. *Logic Journal of the IGPL* 7, 1, 55–78.
- DE ALFARO, L. AND MANNA, Z. 1995. Verification in continuous time by discrete reasoning. In *Proceedings of the 4th International Conference on Algebraic Methodology and Software Technology (AMAST'95)*, V. S. Alagar and M. Nivat, Eds. Lecture Notes in Computer Science, vol. 936. Springer-Verlag, 292–306.
- DE WULF, M., DOYEN, L., MAQUET, N., AND RASKIN, J.-F. 2009. ALASKA: Antichains for Logic, Automata and Symbolic Kripke structures Analysis. In *Proceeding of the 6th International Symposium on Automated Technology for Verification and Analysis (ATVA'08)*. Lecture Notes in Computer Science, vol. 5311. Springer-Verlag, 240–245.
- DEMRI, S. AND SCHNOEBELEN, P. 2002. The complexity of propositional linear temporal logics in simple cases. *Information and Computation* 174, 1, 84–103.
- D'SOUZA, D., MOHAN M., R., AND PRABHAKAR, P. 2007. Flattening metric temporal logic. Manuscript.
- ETESSAMI, K. AND WILKE, T. 1996. An until hierarchy for temporal logic. In *Proceedings of the 11th Annual IEEE Symposium on Logic in Computer Science (LICS'96)*. IEEE Computer Society Press, 108–117.
- FAINEKOS, G. E. AND PAPPAS, G. J. 2007. Robust sampling for MITL specifications. In *Proceedings of the 5th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'07)*. Lecture Notes in Computer Science, vol. 4763. Springer-Verlag, 147–162.
- FAINEKOS, G. E. AND PAPPAS, G. J. 2009. Robustness of temporal logic specifications for continuous-time signals. *Theoretical Computer Science* 410, 42, 4262–4291.
- FURIA, C. A. 2007. Scaling up the formal analysis of real-time systems. Ph.D. thesis, Dipartimento di Elettronica e Informazione, Politecnico di Milano.
- FURIA, C. A., MANDRIOLI, D., MORZENTI, A., AND ROSSI, M. 2010. Modeling time in computing: a taxonomy and a comparative survey. *ACM Computing Surveys* 42, 2 (February), 1–59. Article 6. Also available as <http://arxiv.org/abs/0807.4132>.
- FURIA, C. A., PRADELLA, M., AND ROSSI, M. 2008a. Automated verification of dense-time MTL specifications via discrete-time approximation. In *Proceedings of the 15th International Symposium on Formal Methods (FM'08)*, J. Cuéllar and T. Maibaum, Eds. Lecture Notes in Computer Science, vol. 5014. Springer-Verlag, 132–147.
- FURIA, C. A., PRADELLA, M., AND ROSSI, M. 2008b. Practical automated partial verification of multi-paradigm real-time models. In *Proceedings of the 10th International Conference on Formal Engineering Methods (ICFEM'08)*, S. Liu, T. Maibaum, and K. Araki, Eds. Lecture Notes in Computer Science, vol. 5256. Springer-Verlag, 298–317.
- FURIA, C. A. AND ROSSI, M. 2006. Integrating discrete- and continuous-time metric temporal logics through sampling. In *Proceedings of the 4th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'06)*, E. Asarin and P. Bouyer, Eds. Lecture Notes in Computer Science, vol. 4202. Springer-Verlag, 215–229.
- FURIA, C. A. AND ROSSI, M. 2007. On the expressiveness of MTL variants over dense time. In *Proceedings of the 5th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'07)*, J.-F. Raskin and P. S. Thiagarajan, Eds. Lecture Notes in Computer Science, vol. 4763. Springer-Verlag, 163–178.
- FURIA, C. A. AND ROSSI, M. 2008. MTL with bounded variability: Decidability and complexity. In *Proceedings of the 6th International Conference on Formal Modelling and Analysis of Timed Systems (FORMATS'08)*, F. Cassez and C. Jard, Eds. Lecture Notes in Computer Science, vol. 5215. Springer-Verlag, 109–123.

- GÖLLÜ, A., PURI, A., AND VARAIYA, P. 1994. Discretization of timed automata. In *Proceedings of the 33rd Conference on Decision and Control*. 957–958.
- GRAHAM, R. L., KNUTH, D. E., AND PATASHNIK, O. 1994. *Concrete Mathematics: A foundation for computer science*, 2nd ed. Addison-Wesley.
- HEITMEIER, C. AND MANDRIOLI, D., Eds. 1996. *Formal Methods for Real-Time Computing*. John Wiley & Sons.
- HENZINGER, T. A., MANNA, Z., AND PNUELI, A. 1992. What good are digital clocks? In *Proceedings of the 19th International Colloquium on Automata, Languages and Programming (ICALP'92)*, W. Kuich, Ed. Lecture Notes in Computer Science, vol. 623. Springer-Verlag, 545–558.
- HENZINGER, T. A. AND SIFAKIS, J. 2006. The embedded systems design challenge. In *Proceedings of the 14th International Symposium on Formal Methods (FM'06)*, J. Misra, T. Nipkow, and E. Sekerinski, Eds. Lecture Notes in Computer Science, vol. 4085. Springer-Verlag, 1–15.
- HIRSHFELD, Y. AND RABINOVICH, A. M. 2004. Logics for real time: Decidability and complexity. *Fundamenta Informaticae* 62, 1, 1–28.
- HIRSHFELD, Y. AND RABINOVICH, A. M. 2005. Timer formulas and decidable metric temporal logic. *Information and Computation* 198, 2, 148–178.
- HUNG, D. V. AND GIANG, P. H. 1996. Sampling semantics of duration calculus. In *Proceedings of the 4th International Symposium on Formal Techniques in Real-Time and Fault-Tolerant Systems (FTRTFT'96)*, B. Jonsson and J. Parrow, Eds. Lecture Notes in Computer Science, vol. 1135. Springer-Verlag, 188–207.
- KOYMANS, R. 1990. Specifying real-time properties with metric temporal logic. *Real-Time Systems* 2, 4, 255–299.
- KOYMANS, R. 1992. (real) time: A philosophical perspective. In *Proceedings of the REX Workshop: "Real-Time: Theory in Practice"*, J. W. de Bakker, C. Huizing, W. P. de Roever, and G. Rozenberg, Eds. Lecture Notes in Computer Science, vol. 600. Springer-Verlag, 353–370.
- KRČÁL, P. AND PELÁNEK, R. 2005. On sampled semantics of timed systems. In *Proceedings of the 25th International Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'05)*, R. Ramanujam and S. Sen, Eds. Lecture Notes in Computer Science, vol. 3821. Springer-Verlag, 310–321.
- KUČERA, A. AND STREJČEK, J. 2005. The stuttering principle revisited. *Acta Informatica* 41, 7/8, 415–434.
- MALER, O., NICKOVIC, D., AND PNUELI, A. 2006. From MITL to timed automata. In *Proceedings of the 4th International Conference on Formal Modeling and Analysis of Timed Systems (FORMATS'06)*, E. Asarin and P. Bouyer, Eds. Lecture Notes in Computer Science, vol. 4202. Springer-Verlag, 274–289.
- MALER, O. AND PNUELI, A. 1995. Timing analysis of asynchronous circuits using timed automata. In *Proceedings of the Advanced Research Working Conference on Correct Hardware Design and Verification Methods*, P. Camurati and H. Ekeking, Eds. Lecture Notes in Computer Science, vol. 987. Springer-Verlag, 189–205.
- OUAKNINE, J. 2002. Digitisation and full abstraction for dense-time model checking. In *Proceedings of the 8th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'02)*, J.-P. Katoen and P. Stevens, Eds. Lecture Notes in Computer Science, vol. 2280. Springer-Verlag, 37–51.
- OUAKNINE, J. AND WORRELL, J. 2003. Revisiting digitization, robustness, and decidability for timed automata. In *Proceedings of the 18th Annual IEEE Symposium on Logic in Computer Science (LICS'03)*. IEEE Computer Society Press, 198–207.
- PANDYA, P. K. 2008. Personal communication.
- PANDYA, P. K., NARAYANAN KRISHNA, S., AND LOYA, K. 2007. On sampling abstraction of continuous time logic with durations. In *Proceeding of the 13th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'07)*. Lecture Notes in Computer Science, vol. 4424. Springer-Verlag, 246–260.
- PERRIN, D. AND PIN, J.-É. 2004. *Infinite Words*. Pure and Applied Mathematics, vol. 141. Elsevier.

- PRADELLA, M., MORZENTI, A., AND SAN PIETRO, P. 2007. The symmetry of the past and of the future: Bi-infinite time in the verification of temporal properties. In *Proceedings of The 6th joint meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering (ESEC/FSE 2007)*. 312–320.
- PRADELLA, M., SAN PIETRO, P., SPOLETINI, P., AND MORZENTI, A. 2003. Practical model checking of LTL with past. In *Proceedings of 1st International Workshop on Automated Technology for Verification and Analysis (ATVA'03)*, F. Wang and I. Lee, Eds. Taipei, Taiwan, R.O.C., 135–146.
- RABINOVICH, A. M. 2003. Automata over continuous time. *Theoretical Computer Science* 300, 1–3, 331–363.
- THÉRIEN, D. AND WILKE, T. 2004. Nesting until and since in linear temporal logic. *Theory of Computing Systems* 37, 1, 111–131.

Received December 2009; revised March 2010; accepted April 2010.