Robustness of Temporal Logic Specifications

Georgios E. Fainekos[1] and George J. Pappas[2]

[1] Department of Computer and Information Science, Univ. of Pennsylvania
fainekos@cis.upenn.edu
[2] Department of Electrical and Systems Engineering, Univ. of Pennsylvania
pappasg@ee.upenn.edu

Abstract. In this paper, we consider the robust interpretation of metric temporal logic (MTL) formulas over timed sequences of states. For systems whose states are equipped with nontrivial metrics, such as continuous, hybrid, or general metric transition systems, robustness is not only natural, but also a critical measure of system performance. In this paper, we define robust, multi-valued semantics for MTL formulas, which capture not only the usual Boolean satisfiability of the formula, but also topological information regarding the distance, ε, from unsatisfiability. We prove that any other timed trace which remains ε-close to the initial one also satisfies the same MTL specification with the usual Boolean semantics. We derive a computational procedure for determining an under-approximation to the robustness degree ε of the specification with respect to a given finite timed state sequence. Our approach can be used for robust system simulation and testing, as well as form the basis for simulation-based verification.

Key words: Robustness, Metric spaces, Monitoring, Timed State Sequences, Metric and Linear Temporal Logic

1 Introduction

Model checking [1] has been proven to be a very useful tool for the verification of the properties of software and hardware systems. The tools and methodologies developed for such systems do not naturally extend to systems whose state space is some general metric space, for example linear, nonlinear and hybrid systems. In this case, the model checking problem becomes harder and in most of the cases is undecidable [2]. Therefore, the verification of such systems still relies heavily on methods that involve monitoring and testing [3–6]. Furthermore, general metric transition systems either model physical processes or the interaction between some software and/or hardware system and the continuous physical world. Up to now no formal model exists that can capture accurately the behaviour of such a system – especially if it also exhibits a chaotic behaviour. Moreover, these types of systems have a certain degree of sensitivity with respect to initial conditions or to system parameters. This has one major implication. Deciding the Boolean truth value of a temporal logic specification with respect to a system's trajectory - in some of the cases - does not allow us to draw any conclusions about the real

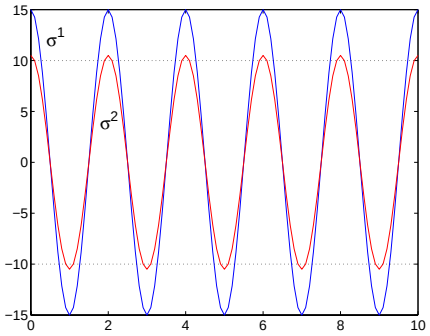

Fig. 1. Two trajectories σ^1 and σ^2 which satisfy the specification: $\Box(\pi_1 \rightarrow \Diamond_{\leq 2}\pi_2)$. Here, $\mathcal{O}(\pi_1) = \mathbb{R}_{\leq -10}$ and $\mathcal{O}(\pi_1) = \mathbb{R}_{\geq 10}$.

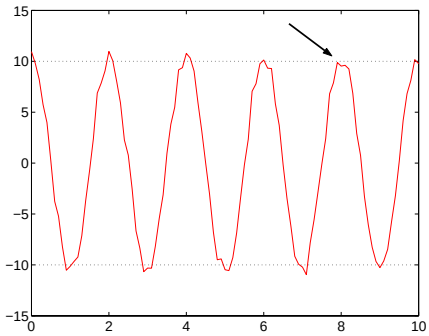

Fig. 2. The trajectory σ^2 modified by random noise. The arrow points to the point in time where the property fails.

system. A small perturbation of the trajectory or the parameters of the system can lead to a different truth value for the formula.

For example, consider the trajectories σ^1 and σ^2 in Fig. 1. Both of them satisfy the same specification "if the value of the state drops below -10, then it should also raise above 10 within 2 time units". Nevertheless, a visual inspection of Fig. 1 indicates that there exists a qualitative difference between σ^1 and σ^2. The later "barely" satisfies the specification. Indeed as we can see in Fig. 2, adding a bounded noise on σ^2 renders the property unsatisfiable on σ^2.

In order to differentiate between such trajectories of a system, we introduce the concept of robustness degree. Informally, we define the robustness degree to be the bound on the perturbation that the trajectory[3] can tolerate without changing the truth value of a specification expressed in the Linear [7] or Metric Temporal Logic [8]. To formally define the robustness degree, we take a topological perspective. We consider finite timed state sequences which take values in some space X equipped with a metric d. If these trajectories are of length n, then each sequence of states is isomorphic to a point in X^n, which is the space of all possible trajectories of length n. In order to quantify how close are two different state sequences in X^n, we define the notion of distance using a metric ρ on the space X^n. Given an LTL or MTL formula ϕ, we can partition the space X^n into two sets: the set P^ϕ of state sequences that satisfy ϕ and the set N^ϕ of state sequences that do not satisfy ϕ. Then, the formal definition of robustness comes naturally, it is just the distance of a state sequence σ from the set P^ϕ or its complement N^ϕ. Using the degree of robustness and the metric ρ, we can

[3] We should bring to notice that we are not interested in the properties of the (possibly) continuous trajectory, but in the properties of its finite representation. Here, we model the finite representation of a continuous trajectory using timed state sequences. Under certain assumptions about the structure of the system, the results in this paper could be mapped back to the continuous case.

define an open ball (tube) around σ and, therefore, we can be sure that any state sequence σ' that remains within the open ball also stays either in P^ϕ or in N^ϕ.

However, the computation of the set P^ϕ and, hence, the computation of the robustness degree are hard problems. To address them, we develop an algorithm that computes an under-approximation of the robustness degree. For that purpose, we define robust semantics for MTL by borrowing ideas from the quantitative version of the linear temporal logic QLTL [9]. Our definition is similar to QLTL (we do not consider discounting), but now the truth values of the MTL formulas range over the closure of the reals instead of the closed interval $[0, 1]$. The atomic propositions in the robust version of MTL evaluate to the distance from the current state in the timed state sequence to the subset of X that the atomic proposition represents. As established in the aforementioned work, the conjunction and disjunction in the Boolean logic are replaced by the min and max operations. Here, the logical negation is replaced by the usual negation of the reals. We prove that when an MTL formula is evaluated with robust semantics over a timed state sequence $\mathcal{T}_1$, then it returns an under-approximation ε of the robustness degree and, therefore, any other timed state sequence $\mathcal{T}_2$ that remains ε-close to $\mathcal{T}_1$ satisfies the same specification. We conclude the paper by presenting a monitoring algorithm (similar to [10, 11]) that is based on the robust semantics of MTL and computes the under-approximation of the robustness degree.

Application-wise the importance of the main contribution of this paper is straightforward: if a system has the property that under bounded disturbances its trajectories remain δ close to the nominal one and, also, its robustness degree with respect to an MTL formula ϕ is $\varepsilon > \delta$, then we know that all the system's trajectories also satisfy the same specification. The timing bounds on the temporal operators, that is the use of MTL instead of LTL, can be justified if one considers that the applications of such a framework are within the systems area. For example, signal processing and simulations of physical systems most of the times do require such constraints. The methodology that we present in this paper can be readily used in several applications such as Qualitative Simulation [12], verification using simulation [13], mobile robot path planning [14] and in behavioral robotics [15].

2 Metric Temporal Logic over Timed State Sequences

2.1 Metric Spaces

Let $\mathbb{R}$ be the set of the real numbers, $\mathbb{Q}$ the set of the rational numbers and $\mathbb{N}$ the set of the natural numbers. We denote the extended real number line by $\overline{\mathbb{R}} = \mathbb{R} \cup \{\pm\infty\}$. Furthermore, we let $\mathbb{B} = \{\top, \bot\}$, where $\top$ and $\bot$ are the symbols for the boolean constants *true* and *false* respectively. If $(X, \leq)$ is a totally ordered set with an ordering relation $\leq$, then an interval of X is denoted by $[a, b]_X = \{x \in X \mid a \leq x \leq b\}$. When $X = \mathbb{R}$, we drop the subscript $\mathbb{R}$. In addition, we use pseudo-arithmetic expressions to represent certain subsets of the aforementioned sets. For example, $\mathbb{R}_{\geq 0}$ denotes the subset of the reals whose

elements are greater or equal to zero. If C is a set, then $cl(C)$ denotes the *closure* of the set C. Let (X, d) be a metric space, i.e. a set X whose topology is induced by the metric d.

Definition 1 (Metric). *A metric on a set X is a positive function $d : X \times X \to \mathbb{R}_{\geq 0}$, such that the three following properties hold*

1. $\forall x_1, x_2, x_3 \in X. d(x_1, x_3) \leq d(x_1, x_2) + d(x_2, x_3)$
2. $\forall x_1, x_2 \in X. d(x_1, x_2) = 0 \Leftrightarrow x_1 = x_2$
3. $\forall x_1, x_2 \in X. d(x_1, x_2) = d(x_2, x_1)$

Using a metric d, we can define the distance of a point $x \in X$ from a set $C \subseteq X$. Intuitively, this distance is the shortest distance from x to all the points in C. In a similar way, the depth of a point x in a set C is defined to be the shortest distance of x from the boundary of C. Both the notions of distance and depth (Fig. 3) will play a fundamental role in the definition of the robustness degree (see Sect. 3).

Definition 2 (Distance, Depth, Signed Distance [16] §8). *Let $x \in X$ be a point, $C \subseteq X$ be a set and d be a metric. Then, we define the*

- Distance *from x to C to be* $\mathbf{dist}_d(x, C) := \inf\{d(x, y) \mid y \in cl(C)\}$
- Depth *of x in C to be* $\mathbf{depth}_d(x, C) := \mathbf{dist}_d(x, X \backslash C)$
- Signed Distance *from x to C to be*

$$\mathbf{Dist}_d(x, C) := \begin{cases} -\mathbf{dist}_d(x, C) & \text{if } x \notin C \\ \mathbf{depth}_d(x, C) & \text{if } x \in C \end{cases}$$

We should point out that we use the extended definition of supremum and infimum, where $\sup \emptyset = -\infty$ and $\inf \emptyset = +\infty$. Also of importance is the notion of an open ball of radius ε centered at a point $x \in X$.

Definition 3 (ε-Ball). *Given a metric d, a radius $\varepsilon \in \overline{\mathbb{R}}_{>0}$ and a point $x \in X$, the open ε-ball centered at x is defined as $B_d(x, \varepsilon) = \{y \in X \mid d(x, y) < \varepsilon\}$.*

It is easy to verify that if the distance ($\mathbf{dist}_d$) of a point x from a set C is $\varepsilon > 0$, then $B_d(x, \varepsilon) \cap C = \emptyset$. And similarly, if $\mathbf{depth}_d(x, C) = \varepsilon > 0$, then $B_d(x, \varepsilon) \subseteq C$.

2.2 Timed State Sequences in Metric Spaces

In this paper, we use *timed state sequences* (TSS) to describe the behavior of a real-time system. Typical models of real time systems are the formalisms of hybrid automata, timed automata, linear and non-linear systems. A *state* of such a system is a point x in a metric space $\mathcal{X} = (X, d)$. With each state of the system x we associate a *time period* Δt, which represents the duration between the occurrence of the current and the previous system states.

Let AP be a finite set of atomic propositions, then the *predicate mapping* $\mathcal{O} : AP \to 2^X$ is a set valued function that assigns to each atomic proposition $\pi \in AP$

a set of states $\mathcal{O}(\pi) \subseteq X$. Furthermore, if the collection of sets $\{\mathcal{O}(\pi)\}_{\pi \in AP}$ is not a cover of X, i.e. $\cup_{\pi \in AP} \mathcal{O}(\pi) \neq X$, then we add to AP a special proposition π_c that maps to the set $\mathcal{O}(\pi_c) = X \backslash \cup_{\pi \in AP} \mathcal{O}(\pi)$. Therefore, we can now define the "inverse" map of $\mathcal{O}$ as $\mathcal{O}^{-1}(x) = \{\pi \in AP \mid x \in \mathcal{O}(\pi)\}$ for $x \in X$. If $x \in \mathcal{O}(\pi)$, then we say that x is a π state. Notice that using the notion of distance, we can quantify how close is a state x to becoming a π state.

The execution of a system can result in an infinite or finite sequence of states. In this paper, we focus on finite sequences of states, which can model the finite representation of a real valued signal or the result of the numerical integration of differential equations.

Definition 4 (TSS). *A timed state sequence $\mathcal{T}$ is a tuple $(\sigma, \tau, \mathcal{O})$ where: $\sigma = x_0, x_1, \ldots, x_n$ is a sequence of states, $\tau = \Delta t_0, \Delta t_1, \ldots, \Delta t_n$ is a sequence of time periods and $\mathcal{O} : AP \to 2^X$ is a predicate mapping; such that $n \in \mathbb{N}$, $x_i \in X$ and $\Delta t_i \in \mathbb{R}_{\geq 0}$ for all $i \in \{0, 1, \ldots, n\}$ and $\Delta t_0, \Delta t_0 + \Delta t_1, \ldots, \sum_{i=0}^{n} \Delta t_i$ is a strictly monotonically increasing sequence.*

We let σ_i and τ_i denote x_i and Δt_i respectively. By convention, we set $\Delta t_0 = 0$. We define $\sigma{\downarrow_i}$ to be the prefix of the state sequence σ, i.e. $\sigma{\downarrow_i} = x_0, x_1, \ldots, x_i$, while $\sigma{\uparrow_i}$ is the suffix, i.e. $\sigma{\uparrow_i} = x_i, x_{i+1}, \ldots, x_n$. The length of $\sigma = x_0, x_1, \ldots, x_n$ is defined to be $|\sigma| = n + 1$. For convenience, we let $|\mathcal{T}| = |\tau| = |\sigma|$ and $\mathcal{T}{\uparrow_i} = (\sigma{\uparrow_i}, \tau{\uparrow_i}, \mathcal{O})$ (similarly for$\downarrow$).

In the following, we use the convention that $\mathcal{T}$ and $\mathcal{S}$ denote the timed state sequences $\mathcal{T} = (\sigma, \tau, \mathcal{O})$ and $\mathcal{S} = (\sigma', \tau, \mathcal{O})$ (and similarly for their superscripted versions). We define Σ_X to be the set of all possible timed state sequences in the space $\mathcal{X} = (X, d)$ and $\Sigma(\mathcal{T})$ to be the set of all possible timed state sequences with the same predicate mapping $\mathcal{O}$ and the same sequence of time periods as $\mathcal{T}$. That is $\Sigma(\mathcal{T}) = \{(\sigma', \tau, \mathcal{O}) \mid \sigma' \in X^{|\mathcal{T}|}\}$. Notice that the sequence σ is isomorphic to a point in the product space $X^{|\sigma|}$.

2.3 Metric Temporal Logic over Finite Timed State Sequences

The Metric Temporal Logic (MTL) [8] is an extension of the Linear Temporal Logic (LTL) [7]. In MTL, the syntax of the logic is extended to include timing constraints on the usual temporal operators of LTL. Using LTL specifications we can check qualitative timing properties, while with MTL specifications quantitative timing properties. Recently, it was shown by Ouaknine and Worrell [17] that MTL is decidable over finite timed state sequences. In this section, we review the basics of MTL with point-based semantics (as opposed to interval based semantics [18]) over finite timed state sequences.

Definition 5 (Syntax of MTL). *Let AP be the set of atomic propositions, D the set of truth degree constants and $\mathcal{I}$ an interval of $\mathbb{R}_{\geq 0}$ with rational endpoints. The set Φ_D of all well-formed formulas (wff) is the smallest set such that*

- *it contains all the members of D and AP, i.e. $D, AP \subseteq \Phi_D$*
- *if $\phi_1, \phi_2 \in \Phi_D$, then $\neg\phi_1, \phi_1 \vee \phi_2, \bigcirc_{\mathcal{I}}\phi_1, \phi_1 \mathcal{U}_{\mathcal{I}} \phi_2$ belong to Φ_D*

In the following, we fix the set AP, while the set D varies. As usual, $\phi_1 \wedge \phi_2 = \neg(\neg\phi_1 \vee \neg\phi_2)$ and $\phi_1 \rightarrow \phi_2 = \neg\phi_1 \vee \phi_2$. Here, $\bigcirc_{\mathcal{I}}$ is the *next* time operator and $\mathcal{U}_{\mathcal{I}}$ the *until* operator. We can also define the common temporal operators *eventually* $\Diamond_{\mathcal{I}}\phi = \top\,\mathcal{U}_{\mathcal{I}}\phi$ and *always* $\Box_{\mathcal{I}}\phi = \neg\Diamond_{\mathcal{I}}\neg\phi$. In the case where $\mathcal{I} = [0, +\infty)$, we remove the subscript $\mathcal{I}$ from the temporal operators, i.e. we just write $\mathcal{U}$, $\bigcirc$, $\Diamond$ and $\Box$. When all the subscripts of the temporal operators are of the form $[0, +\infty)$, then the MTL formula ϕ reduces to an LTL formula and we can ignore the time periods.

The subscript $\mathcal{I}$ imposes timing constraints on the temporal operators. The interval $\mathcal{I}$ can be open, half-open or closed, bounded or unbounded. The function lb returns the lower (or left) bound of the interval $\mathcal{I}$ whereas the function ub returns the upper (or right) bound. Note that $lb(\mathcal{I}), ub(\mathcal{I}) \in \mathbb{Q}_{\geq 0}$ and that it could be the case that $ub(\mathcal{I}) = lb(\mathcal{I})$, i.e. $\mathcal{I}$ is a singleton. For any $t \in \mathbb{Q}$, we define $\mathcal{I} + t = \{t' + t \mid t' \in \mathcal{I}\}$. Also, we do not consider relative [10] and absolute congruences [19] and we have not included the *since* and *last* temporal operators (the past fragment) in the syntax of MTL.

Metric Temporal Logic (MTL) formulas are interpreted over timed state sequences $\mathcal{T}$ with $|\mathcal{T}| > 0$. The constraint $|\mathcal{T}| > 0$ implies that the sequence has at least one state, that is we ignore the pathological cases of empty state sequences. In this paper, we denote formula satisfiability using a membership function $\langle\!\langle \phi \rangle\!\rangle : \Sigma_X \rightarrow \mathbb{B}$ instead of the usual notation $\mathcal{T} \models \phi$. The functional approach enables us to maintain a uniform presentation throughout this paper. We say that a timed state sequence $\mathcal{T}$ satisfies the formula ϕ when $\langle\!\langle \phi \rangle\!\rangle(\mathcal{T}) = \top$. In this case, we refer to $\mathcal{T}$ as a *model* of ϕ. The set of all models of ϕ is denoted by $\mathcal{L}(\phi)$, i.e. $\mathcal{L}(\phi) = \{\mathcal{T} \in \Sigma_X \mid \langle\!\langle \phi \rangle\!\rangle(\mathcal{T}) = \top\}$.

Definition 6 (Semantics of MTL). *Let $\mathcal{T} = (\sigma, \tau, \mathcal{O}) \in \Sigma_X$, $v \in \mathbb{B}$, $\pi \in AP$, $i, j \in \mathbb{N}$ and $K_{\mathcal{I}}^{\mathcal{T}} = \{i \in [0, |\mathcal{T}| - 1]_{\mathbb{N}} \mid \sum_{j=0}^{i} \tau_j \in \mathcal{I}\}$, then the semantics[4] of any formula $\phi \in \Phi_{\mathbb{B}}$ are inductively defined by*

$$
\begin{aligned}
\langle\!\langle v \rangle\!\rangle(\mathcal{T}) &:= v \\
\langle\!\langle \pi \rangle\!\rangle(\mathcal{T}) &:= \sigma_0 \in \mathcal{O}(\pi) \\
\langle\!\langle \neg\psi \rangle\!\rangle(\mathcal{T}) &:= \neg\langle\!\langle \psi \rangle\!\rangle(\mathcal{T}) \\
\langle\!\langle \phi_1 \vee \phi_2 \rangle\!\rangle(\mathcal{T}) &:= \langle\!\langle \phi_1 \rangle\!\rangle(\mathcal{T}) \vee \langle\!\langle \phi_2 \rangle\!\rangle(\mathcal{T}) \\
\langle\!\langle \bigcirc_{\mathcal{I}}\psi \rangle\!\rangle(\mathcal{T}) &:= \begin{cases} (\tau_1 \in \mathcal{I}) \wedge \langle\!\langle \psi \rangle\!\rangle(\mathcal{T}\!\uparrow_1) & \text{if } |\mathcal{T}| > 1 \\ \bot & \text{otherwise} \end{cases} \\
\langle\!\langle \phi_1\,\mathcal{U}_{\mathcal{I}}\phi_2 \rangle\!\rangle(\mathcal{T}) &:= \bigvee_{i=0}^{|\mathcal{T}|-1} \big((i \in K_{\mathcal{I}}^{\mathcal{T}}) \wedge \langle\!\langle \phi_2 \rangle\!\rangle(\mathcal{T}\!\uparrow_i) \wedge \bigwedge_{j=0}^{i-1} \langle\!\langle \phi_1 \rangle\!\rangle(\mathcal{T}\!\uparrow_j)\big)
\end{aligned}
$$

Informally, the path formula $\phi_1\,\mathcal{U}_{[a,b]}\phi_2$ expresses the property that over the timed state sequence $\mathcal{T}$ and in the time interval $[a, b]$, ϕ_2 becomes true and for all previous time ϕ_1 holds.

[4] Note that here we overload the symbols and we use the same notation for both the logical connectives in the MTL formulas and their respective Boolean truth degree functions.

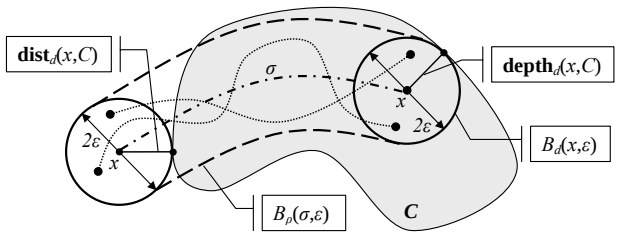

Fig. 3. A tube (*dashed lines*) around a nominal state sequence σ (*dash-dotted line*). The tube encloses a set of state sequences (*dotted lines*). Also, the definition of distance and depth and the associated neighborhoods.

3 Robust Satisfaction of MTL Specifications

3.1 Toward a Notion of Robust Satisfaction

In this section, we define what it means for a timed state sequence (taking values in some metric space) to satisfy a Metric Temporal Logic specification *robustly*. In the case of the timed state sequences that we consider in this paper, we can quantify how close are two different state sequences by using the metric d. Let $\mathcal{T} = (\sigma, \tau, \mathcal{O})$ be a timed state sequence and $(\sigma', \tau, \mathcal{O}) \in \Sigma(\mathcal{T})$, then

$$\rho(\sigma, \sigma') = \max\{d(\sigma_i, \sigma'_i) \mid i \in [0, |\sigma| - 1]_{\mathbb{N}}\} \tag{1}$$

is a metric on the set $X^{|\mathcal{T}|}$, which is well defined since $|\mathcal{T}|$ is finite. Now that the space of state sequences is equipped with a metric, we can define a tube around a timed state sequence $\mathcal{T}$. Given an $\varepsilon > 0$, we let

$$\Sigma_\varepsilon(\mathcal{T}) = \{(\sigma', \tau, \mathcal{O}) \in \Sigma(\mathcal{T}) \mid \sigma' \in B_\rho(\sigma, \varepsilon)\}$$

to be the set of all timed state sequences that remain ε-close to $\mathcal{T}$.

Informally, we define the degree of robustness that a timed state sequence $\mathcal{T}$ satisfies an MTL formula ϕ to be a number $\varepsilon \in \overline{\mathbb{R}}$. Intuitively, a positive ε means that the formula ϕ is satisfiable and, moreover, that all the other timed state sequences that remain ε-close to the nominal one also satisfy ϕ. Accordingly, if ε is negative, then $\mathcal{T}$ does not satisfy ϕ and all the other timed state sequences that remain within the open tube of radius $|\varepsilon|$ also do not satisfy ϕ.

Definition 7 (Robustness Degree). *Let $\phi \in \Phi_{\mathbb{B}}$, $\mathcal{T} = (\sigma, \tau, \mathcal{O}) \in \Sigma_X$ and ρ be the metric (1). Define $P^\phi_\mathcal{T} := \{\sigma' \mid (\sigma', \tau, \mathcal{O}) \in \Sigma(\mathcal{T}) \cap \mathcal{L}(\phi)\}$, then the robustness degree $\varepsilon \in \overline{\mathbb{R}}$ of $\mathcal{T}$ with respect to ϕ is defined as $\varepsilon := \mathbf{Dist}_\rho(\sigma, P^\phi_\mathcal{T})$.*

Remark 1. $P^\phi_\mathcal{T}$ is the set of all models with a sequence of time periods τ that satisfy ϕ. If we define $N^\phi_\mathcal{T} := \{\sigma' \mid (\sigma', \tau, \mathcal{O}) \in \Sigma(\mathcal{T}) \cap \Sigma_X \backslash \mathcal{L}(\phi)\}$, then the set $\{P^\phi_\mathcal{T}, N^\phi_\mathcal{T}\}$ forms a partition of the set $X^{|\mathcal{T}|}$. Therefore, we have duality $P^\phi_\mathcal{T} = X^{|\mathcal{T}|} \backslash N^\phi_\mathcal{T}$ and $N^\phi_\mathcal{T} = X^{|\mathcal{T}|} \backslash P^\phi_\mathcal{T}$.

The following proposition is derived directly from the definitions. It states that all the timed state sequences $\mathcal{S}$, which have distance from $\mathcal{T}$ less than the robustness degree of $\mathcal{T}$ with respect to ϕ, satisfy the same specification ϕ as $\mathcal{T}$.

Proposition 1. *Let $\phi \in \Phi_{\mathbb{B}}$, $\mathcal{T} = (\sigma, \tau, \mathcal{O}) \in \Sigma_X$ and $\varepsilon = \mathbf{Dist}_\rho(\sigma, P^\phi_{\mathcal{T}})$. If $|\varepsilon| > 0$, then for all $\mathcal{S} \in \Sigma_{|\varepsilon|}(\mathcal{T})$ it is $\langle\!\langle \phi \rangle\!\rangle(\mathcal{S}) = \langle\!\langle \phi \rangle\!\rangle(\mathcal{T})$.*

Remark 2. If $\varepsilon = 0$, then the truth value of ϕ with respect to $\mathcal{T}$ is not robust, i.e. any small perturbation of a critical state in the timed state sequence can change the satisfiability of the formula with respect to $\mathcal{T}$.

Theoretically, the set $P^\phi_{\mathcal{T}}$ (or $N^\phi_{\mathcal{T}}$) can be computed. A naive, but straightforward, way to construct the set $P^\phi_{\mathcal{T}}$ is as follows. Instead of timed state sequences in a metric space X, let us consider finite timed state sequences where each state is a set of atomic propositions. We will refer to the later as timed words for clarity. In more detail, consider the timed word $\mathcal{T}_w = (\xi, \tau)$ where for all $i = 0, 1, \ldots, |\mathcal{T}_w| - 1$ it is $\xi_i \in \overline{AP} = 2^{AP} \backslash \emptyset$. In [17], it was proven the one can construct an acceptor $\mathcal{A}_\phi$ (in the form of a timed alternating automaton with one clock) for the finite models $\mathcal{T}_w$ of any formula ϕ in the logic MTL with the standard semantics (that is $\langle\!\langle \pi \rangle\!\rangle(\mathcal{T}_w) := \pi \in \xi_0$). Assume now that we are given an MTL formula ϕ, a sequence of time periods τ and a predicate mapping $\mathcal{O}$. For that particular τ, we can find the set $\mathcal{L}_\tau(\mathcal{A}_\phi)$ of timed words (ξ, τ) that are accepted by $\mathcal{A}_\phi$. One way to do so is to construct the set UW_τ of all possible untimed words ξ of length $|\tau|$, that is $UW_\tau = \overline{AP}^{|\tau|}$, and, then, for each $\xi \in UW_\tau$ verify whether (ξ, τ) is accepted by $\mathcal{A}_\phi$, i.e. whether $(\xi, \tau) \in \mathcal{L}(\mathcal{A}_\phi)$ and, thus, $(\xi, \tau) \in \mathcal{L}_\tau(\mathcal{A}_\phi)$. This can be done in time $O(|\tau||\overline{AP}|^{|\tau|})$ since given the automaton $\mathcal{A}_\phi$ it takes linear time in the length of the timed word to decide whether the word is in the language or not. From the set $\mathcal{L}_\tau(\mathcal{A}_\phi)$, we can easily derive the set $P^\phi_{\mathcal{T}} = \bigcup_{(\xi,\tau)\in\mathcal{L}_\tau(\mathcal{A}_\phi)} \big((\cap_{\pi\in\xi_0}\mathcal{O}(\pi)) \times \ldots \times (\cap_{\pi\in\xi_{|\mathcal{T}|-1}}\mathcal{O}(\pi))\big)$.

The following toy example illustrates the concept of robustness for temporal logic formulas interpreted over finite (timed) state sequences.

Example 1. Assume that we are given the LTL specification $\phi = \pi_1 \mathcal{U} \pi_2$ such that $\mathcal{O}(\pi_1) = [1, 2] \subseteq \mathbb{R}$ and $\mathcal{O}(\pi_2) = [0, 1) \subseteq \mathbb{R}$. Moreover, we have $\mathcal{O}(\pi_c) =$

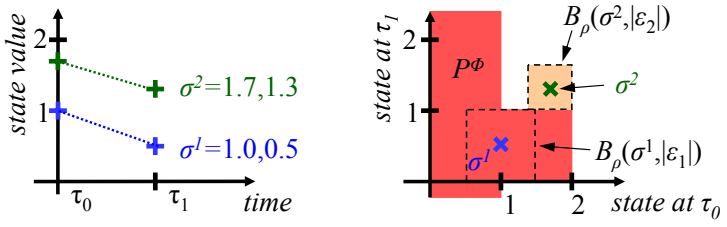

Fig. 4. On the left appears the time-domain representation of the timed state sequences $\mathcal{T}_1$ (*blue crosses*) and $\mathcal{T}_2$ (*green crosses*) of Example 1. On the right appears the space of the state sequences of length 2. Each x represents a state sequence as a point in $\mathbb{R}^2$.

$\mathbb{R}\backslash(\mathcal{O}(\pi_1) \cup \mathcal{O}(\pi_2)) = (-\infty, 0) \cup (2, +\infty)$. Note that the sets $\mathcal{O}(\pi_1)$, $\mathcal{O}(\pi_2)$ and $\mathcal{O}(\pi_c)$ are mutually disjoint. Consider now two timed state sequences $\mathcal{T}_1 = (\sigma^1, \tau, \mathcal{O})$ and $\mathcal{T}_2 = (\sigma^2, \tau, \mathcal{O})$ taking values in $\mathbb{R}$ such that $\sigma^1 = 1, 0.5$ and $\sigma^2 = 1.7, 1.3$. Since ϕ is an LTL formula, we can ignore the sequence of time periods τ. In this simple case, we can compute the set P^{Φ} with the procedure described above. The four untimed words that satisfy the specification ϕ and generate non-empty sets are $\xi^1 = \{\pi_2\}, \{\pi_1\}$, $\xi^2 = \{\pi_2\}, \{\pi_2\}$, $\xi^3 = \{\pi_2\}, \{\pi_c\}$ and $\xi^4 = \{\pi_1\}, \{\pi_2\}$. Therefore, we get $P^{\phi} = P^{\phi}_{\mathcal{T}_1} = P^{\phi}_{\mathcal{T}_2} = \mathcal{O}(\pi_2) \times \mathcal{O}(\pi_1) \cup \mathcal{O}(\pi_2) \times \mathcal{O}(\pi_2) \cup \mathcal{O}(\pi_2) \times \mathcal{O}(\pi_c) \cup \mathcal{O}(\pi_1) \times \mathcal{O}(\pi_2) = [0, 1) \times \mathbb{R} \cup [1, 2] \times [0, 1)$ (see Fig. 4). Therefore, $\varepsilon_1 = \mathbf{Dist}_{\rho}(\sigma^1, P^{\phi}) = 0.5$ and $\varepsilon_2 = \mathbf{Dist}_{\rho}(\sigma^2, P^{\phi}) = -0.3$.

3.2 Computing an Under-Approximation of the Robustness Degree

The aforementioned theoretical construction of the set $P^{\phi}_{\mathcal{T}}$ cannot be of any practical interest. Moreover, the definition of robustness degree involves a number of set operations (union, intersection and complementation) in the possibly high dimensional spaces X and $X^{|\mathcal{T}|}$, which can be computationally expensive in practice. Therefore in this section, we develop an algorithm that computes an under-approximation of the robustness degree ε by directly operating on the timed state sequence while avoiding set operations. In the following, we refer to the approximation of the robustness degree as the *robustness estimate*. As it is usually the case in trade-offs, we gain computational efficiency at the expense of accuracy.

In order to compute the robustness estimate, we define robust semantics for MTL. For this purpose, we extend the classical notion of formula satisfiability to the multi-valued case. In this framework, each formula takes truth values over a finite or infinite set of values that have an associated partial or total order relation. In this paper, we differentiate from previous works [9] by providing the definition of multi-valued semantics for MTL based on robustness considerations.

Let $\mathfrak{R} = (\overline{\mathbb{R}}, \leq)$ be the closure of the reals with the usual ordering relation. We define the binary operators $\sqcup : \overline{\mathbb{R}} \times \overline{\mathbb{R}} \to \overline{\mathbb{R}}$ and $\sqcap : \overline{\mathbb{R}} \times \overline{\mathbb{R}} \to \overline{\mathbb{R}}$ using the maximum and minimum functions as $x \sqcup y := \max\{x, y\}$ and $x \sqcap y := \min\{x, y\}$. Also, for some $R \subseteq \overline{\mathbb{R}}$ we extend the above definitions as follows $\bigsqcup R := \sup R$ and $\bigsqcap R := \inf R$. Recall that $\bigsqcup \overline{\mathbb{R}} = +\infty$ and $\bigsqcap \overline{\mathbb{R}} = -\infty$ and that any subset of $\overline{\mathbb{R}}$ has a supremum and infimum. Finally, because $\mathfrak{R}$ is a totally ordered set, it is *distributive*, i.e. for all $a, b, c \in \overline{\mathbb{R}}$ it is $a \sqcap (b \sqcup c) = (a \sqcap b) \sqcup (a \sqcap c)$ and $a \sqcup (b \sqcap c) = (a \sqcup b) \sqcap (a \sqcup c)$.

We propose multi-valued semantics for the Metric Temporal Logic where the valuation function on the atomic propositions takes values over the totally ordered set $\mathfrak{R}$ according to the metric d operating on the state space X of the timed state sequence $\mathcal{T}$. For this purpose, we let the valuation function to be the signed distance from the current point in the state sequence σ_0 to a set C labeled by the atomic proposition. Intuitively, this distance represents how robustly is the point σ_0 within a set C. If this metric is zero, then even the smallest perturbation of the point can drive it inside or outside the set C, dramatically affecting membership.

For the purposes of the following discussion, we use the notation $[\![\phi]\!](\mathcal{T})$ to denote the robustness estimate with which the structure $\mathcal{T}$ satisfies the specification ϕ (formally $[\![\phi]\!] : \Sigma_X \to \overline{\mathbb{R}}$).

Definition 8 (Robust Semantics of MTL). *Let $\mathcal{T} = (\sigma, \tau, \mathcal{O}) \in \Sigma_X$, $v \in \overline{\mathbb{R}}$, $\pi \in AP$, $i, j \in \mathbb{N}$ and $K^{\mathcal{T}}_{\mathcal{I}} = \{i \in [0, |\mathcal{T}| - 1]_{\mathbb{N}} \mid \sum_{j=0}^{i} \tau_j \in \mathcal{I}\}$, then the robust semantics of a formula $\phi \in \Phi_{\overline{\mathbb{R}}}$ with respect to $\mathcal{T}$ are inductively defined by*

$$
\begin{aligned}
[\![v]\!](\mathcal{T}) &:= v \\
[\![\pi]\!](\mathcal{T}) &:= \mathbf{Dist}_d(\sigma_0, \mathcal{O}(\pi)) \\
[\![\neg\psi]\!](\mathcal{T}) &:= -[\![\psi]\!](\mathcal{T}) \\
[\![\phi_1 \vee \phi_2]\!](\mathcal{T}) &:= [\![\phi_1]\!](\mathcal{T}) \sqcup [\![\phi_2]\!](\mathcal{T}) \\
[\![\bigcirc_{\mathcal{I}}\psi]\!](\mathcal{T}) &:= \begin{cases} \mathbf{mv}(\tau_1 \in \mathcal{I}) \sqcap [\![\psi]\!](\mathcal{T}\!\uparrow_1) & \text{if } |\mathcal{T}| > 1 \\ -\infty & \text{otherwise} \end{cases} \\
[\![\phi_1 \,\mathcal{U}_{\mathcal{I}} \phi_2]\!](\mathcal{T}) &:= \bigsqcup_{i=0}^{|\mathcal{T}|-1} \big(\mathbf{mv}(i \in K^{\mathcal{T}}_{\mathcal{I}}) \sqcap [\![\phi_2]\!](\mathcal{T}\!\uparrow_i) \sqcap \textstyle\bigsqcap_{j=0}^{i-1} [\![\phi_1]\!](\mathcal{T}\!\uparrow_j)\big)
\end{aligned}
$$

where the unary operator $(-)$ is defined to be the negation over the reals.

Remark 3. It is easy to verify that the semantics of the negation operator give us all the usual nice properties such as the *De Morgan laws*: $a \sqcup b = -(-a \sqcap -b)$ and $a \sqcap b = -(-a \sqcup -b)$, *involution*: $-(-a) = a$ and *antisymmetry*: $a \leq b$ iff $-a \geq -b$ for $a, b \in \overline{\mathbb{R}}$.

Since the truth degree constants of the formulas in $\Phi_{\mathbb{B}}$ differ from those of the formulas in $\Phi_{\overline{\mathbb{R}}}$, we define a translation function $\mathbf{mv} : \Phi_{\mathbb{B}} \to \Phi_{\overline{\mathbb{R}}}$ which takes as input a formula $\phi \in \Phi_{\mathbb{B}}$ and replaces the occurrences of $\bot$ and $\top$ by $-\infty$ and $+\infty$ respectively. All the other symbols in ϕ are left the same. The following proposition states the relationship between the usual and the robust semantics of MTL (the proof uses induction on the structure of ϕ).

Proposition 2 (proof in [20]). *Let $\phi \in \Phi_{\mathbb{B}}$, $\psi = \mathbf{mv}(\phi)$ and $\mathcal{T} \in \Sigma_X$, then*

$$
\begin{array}{llll}
(1) & [\![\psi]\!](\mathcal{T}) > 0 \;\Rightarrow\; \langle\!\langle\phi\rangle\!\rangle(\mathcal{T}) = \top & (2) & \langle\!\langle\phi\rangle\!\rangle(\mathcal{T}) = \top \;\Rightarrow\; [\![\psi]\!](\mathcal{T}) \geq 0 \\
(3) & [\![\psi]\!](\mathcal{T}) < 0 \;\Rightarrow\; \langle\!\langle\phi\rangle\!\rangle(\mathcal{T}) = \bot & (4) & \langle\!\langle\phi\rangle\!\rangle(\mathcal{T}) = \bot \;\Rightarrow\; [\![\psi]\!](\mathcal{T}) \leq 0
\end{array}
$$

Note that the equivalence in the above proposition fails because, if a point is on the boundary of the set, its distance to the set or its depth in the set is by definition zero. Therefore, the point is classified to belong to that set even if the set is open in the topology.

The following theorem identifies the robustness estimate as an underapproximation of the robustness degree (proof by induction on the structure of ϕ).

Theorem 1 (proof in [20]). *Given $\phi \in \Phi_{\mathbb{B}}$ and $\mathcal{T} = (\sigma, \tau, \mathcal{O}) \in \Sigma_X$, then*

$$
|[\![\mathbf{mv}(\phi)]\!](\mathcal{T})| \leq |\mathbf{Dist}_{\rho}(\sigma, P^{\phi}_{\mathcal{T}})| \tag{2}
$$

In more detail, $-\mathbf{depth}_{\rho}(\sigma, N^{\phi}_{\mathcal{T}}) \leq [\![\phi]\!](\mathcal{T}) \leq \mathbf{depth}_{\rho}(\sigma, P^{\phi}_{\mathcal{T}})$.

In the above theorem, the equality in equation (2) fails due to the robust interpretation of the disjunction connective. The inequality manifests itself in four distinct ways: (i) at the level of the atomic propositions, i.e. $\pi_1 \vee \pi_2$, (ii) due to the existence of tautologies in the formula, i.e. $\pi \vee \neg\pi$, (iii) when we consider disjuncts of MTL subformulas, i.e. $\phi_1 \vee \phi_2$, and more importantly, (iv) due to the disjunctions in the semantics of the until temporal operator.

The first case can be remedied by introducing a new symbol for each Boolean combination of atomic propositions. The second and third conditions require the attention of the user of the algorithm. Even though the above cases can be fixed by introducing syntactic restrictions, the last case (iv) captures a fundamental shortcoming of the robust semantics. The timed state sequences that have state sequences in $B_\rho(\sigma, |\mathbf{Dist}_\rho(\sigma, P^\phi_{\mathcal{T}})|)$ can satisfy or falsify the specification ϕ at different time instants than $\mathcal{T}$. On the other hand, the robustness estimate returns the "radius" of the neighborhood of traces that satisfy the specification at the same point in time.

Example 2. Going back to Example 1, we have seen that $\varepsilon_1 = \mathbf{Dist}_\rho(\sigma^1, P^\phi) = 0.5$. Nevertheless, $[\![\phi]\!](\mathcal{T}_1) = [\![\pi_2]\!](\mathcal{T}_1) \sqcup ([\![\pi_1]\!](\mathcal{T}_1) \sqcap [\![\pi_2]\!](\mathcal{T}_1\uparrow_1)) = 0 \sqcup (0 \sqcap 0.5) = 0 \neq \varepsilon_1$. Consider now a timed state sequence $\mathcal{T}' = (\sigma', \tau, \mathcal{O})$ such that $\sigma' = 1.1, 0.5$. It is immediate to see that $\langle\!\langle\phi\rangle\!\rangle(\mathcal{T}') = \top$ and that $\mathcal{T}' \in \Sigma_{\varepsilon_1}(\mathcal{T}_1)$. Note that $\mathcal{T}_1$ satisfies the specification at time τ_1, while $\mathcal{T}'$ satisfies ϕ at time τ_0. The robust semantics of MTL cannot capture this.

From Proposition 1 and Theorem 1 we derive the next theorem as a corollary.

Theorem 2. *Given $\phi \in \Phi_{\mathbb{B}}$ and $\mathcal{T} \in \Sigma_X$, if $[\![\mathbf{mv}(\phi)]\!](\mathcal{T}) = \varepsilon$ and $|\varepsilon| > 0$, then for all $\mathcal{S} \in \Sigma_{|\varepsilon|}(\mathcal{T})$ it is $\langle\!\langle\phi\rangle\!\rangle(\mathcal{S}) = \langle\!\langle\phi\rangle\!\rangle(\mathcal{T})$.*

Theorem 2 has several implications. First, in the simplest case where we just simulate the response of a system, we can derive bounds for the magnitude of the disturbances that the system can tolerate while still satisfying the same MTL specification. Second, we can use approximation metrics [21] in order to verify a system using simulations [22].

4 Monitoring the Robustness of Temporal Properties

In this section, we present a procedure that computes the robustness estimate of a timed state sequence $\mathcal{T}$ with respect to a specification ϕ stated in the Metric Temporal Logic. For this purpose, we design a monitoring algorithm based on the classical and robust semantics of MTL.

Starting from the definition of the Boolean semantics of the until operator and using the distributive law, we can derive an equivalent recursive formulation (see also [10]):

$$\langle\!\langle\phi_1 \,\mathcal{U}_{\mathcal{I}} \phi_2\rangle\!\rangle(\mathcal{T}) = \begin{cases} ((0 \in \mathcal{I}) \wedge \langle\!\langle\phi_2\rangle\!\rangle(\mathcal{T})) \vee \\ \vee(\langle\!\langle\phi_1\rangle\!\rangle(\mathcal{T}) \wedge \langle\!\langle\phi_1 \,\mathcal{U}_{\mathcal{I}-\tau_1} \phi_2\rangle\!\rangle(\mathcal{T}\uparrow_1)) & \text{if } |\mathcal{T}| > 1 \\ (0 \in \mathcal{I}) \wedge \langle\!\langle\phi_2\rangle\!\rangle(\mathcal{T}) & \text{otherwise} \end{cases}$$

Algorithm 1 Monitoring the Robustness of Timed State Sequences

Input: The MTL formula ϕ and the timed state sequence $\mathcal{T} = (\sigma, \tau, \mathcal{O})$
Output: The formula's Boolean truth value and the robustness parameter

procedure MONITOR(ϕ,$\mathcal{T}$)
  if $|\mathcal{T}| > 1$ **then return** $\phi \leftarrow$ PROGRESS$(\phi, \sigma_0, \tau_1, \bot, \mathcal{O})$
  else return $\phi \leftarrow$ PROGRESS$(\phi, \sigma_0, 0, \top, \mathcal{O})$
  end if
  if $\phi = (v, \varepsilon)$ **then return** (v, ε) ▷ $v \in \{\top, \bot\}$ and $\varepsilon \in \overline{\mathbb{R}}$
  else return MONITOR$(\phi, \mathcal{T}\uparrow_1)$
  end if
end procedure

A similar recursive formulation holds for the robust MTL semantics (see [20]). Using the recursive definitions, it is easy to derive an algorithm that returns the Boolean truth value[5] of the formula and its robustness degree. The main observation is that each value node in the parse tree of the MTL formula should also contain its robustness degree. Therefore, the only operations that we need to modify are the negation and disjunction which must perform, respectively, a negation and a maximum operation on the robustness values of their operants. Then, the new semantics for the conjunction operator can be easily derived from these two.

Definition 9 (Hybrid Semantics for Negation and Disjunction). *Let* $(v_1, \varepsilon_1), (v_2, \varepsilon_2) \in \mathbb{B} \times \overline{\mathbb{R}}$*, then we define*

- *Negation:* $\neg(v, \varepsilon) := (\neg v, -\varepsilon)$
- *Disjunction:* $(v_1, \varepsilon_1) \vee (v_2, \varepsilon_2) := (v_1 \vee v_2, \max\{\varepsilon_1, \varepsilon_2\})$

Given a timed state sequence $\mathcal{T}$ and an MTL formula ϕ, we can construct a monitoring algorithm (Algorithm 1) that can decide both the satisfaction of the formula and the robustness parameter ε on-the-fly. Algorithm 2 is the core of the monitoring procedure. It takes as input the temporal logic formula ϕ, the current state s and the time period before the next state occurs, it evaluates the part of the formula that must hold on the current state and returns the formula that it has to hold at the next state of the timed trace. In Algorithm 2, $\overleftarrow{\mathcal{I}}$ is defined as follows

$$\overleftarrow{\mathcal{I}} = \begin{cases} [0, lb(\mathcal{I})] \cup \mathcal{I} & \text{if } 0 < lb(\mathcal{I}) \\ \mathcal{I} & \text{otherwise} \end{cases}$$

The constraint $0 \in \overleftarrow{\mathcal{I}}$ is added in order to terminate the propagation of the subformula $\phi_1 \mathcal{U}_{\mathcal{I}-\tau_1} \phi_2$, when the timing constraints for the occurrence of ϕ_2 have already been violated. Note that this timing constraint is meaningful only if we also perform the following simplifications at each recursive call of the algorithm PROGRESS: (i) $\phi \wedge (\top, +\infty) \equiv \phi$, (ii) $\phi \vee (\bot, -\infty) \equiv \phi$, (iii) $\phi \vee (\top, +\infty) \equiv (\top, +\infty)$ and (iv) $\phi \wedge (\bot, -\infty) \equiv (\bot, -\infty)$.

[5] Note that the Boolean truth valued is required in the cases where the robustness degree is zero (see Proposition 2).

Algorithm 2 Formula Progression Algorithm

Input: The MTL formula ϕ, the current state s, the time period Δt for the next state, a variable *last* indicating whether the next state is the last and the mapping $\mathcal{O}$
Output: The MTL formula ϕ that has to hold at the next state

procedure PROGRESS($\phi, s, \Delta t, last, \mathcal{O}$)
  if $\phi = (v, \varepsilon) \in \{\bot, \top\} \times \overline{\mathbb{R}}$ **then return** (v, ε)
  else if $\phi = \pi$ **then return** $(s \in \mathcal{O}(\pi), \mathbf{Dist}_d(s, \mathcal{O}(\pi)))$
  else if $\phi = \neg\psi$ **then return** $\neg$PROGRESS($\psi, s, \Delta t, last, \mathcal{O}$)
  else if $\phi = \phi_1 \vee \phi_2$ **then**
    return PROGRESS($\phi_1, s, \Delta t, last, \mathcal{O}$)$\vee$ PROGRESS($\phi_2, s, \Delta t, last, \mathcal{O}$)
  else if $\phi = \bigcirc_{\mathcal{I}}\psi$ **then return** HYBRID($\neg last \wedge (\Delta t \in \mathcal{I})) \wedge \psi$
  else if $\phi = \phi_1 \mathcal{U}_{\mathcal{I}} \phi_2$ **then**
    $\alpha \leftarrow$ HYBRID($0 \in \mathcal{I}$)$\wedge$ PROGRESS($\phi_2, s, \Delta t, last, \mathcal{O}$)
    $\beta \leftarrow$ HYBRID($\neg last \wedge (0 \in \overleftarrow{\mathcal{I}})$)$\wedge$ PROGRESS($\phi_1, s, \Delta t, last, \mathcal{O}) \wedge \phi_1 \mathcal{U}_{\mathcal{I}-\Delta t}\phi_2$
    return $\alpha \vee \beta$
  end if
end procedure

function HYBRID($Bool$)
  if $Bool = \top$ **return** $(\top, +\infty)$ **else return** $(\bot, -\infty)$ **end if**
end function

When we check how robustly a timed state sequence $\mathcal{T}$ satisfies a specification ϕ, we cannot stop the monitoring process as soon as we can determine that the MTL formula holds on $\mathcal{T}$. This is because a future state in the timed state sequence may satisfy the specification more robustly. Therefore, it is preferable to execute the procedure MONITOR for the whole length of the timed state sequence $\mathcal{T}$.

The proof of the following theorem is standard and uses induction on the structure of ϕ based on the classical and robust semantics of MTL.

Theorem 3 (proof in [20]). *Given an MTL formula $\phi \in \Phi_{\mathbb{B}}$ and a timed state sequence $\mathcal{T} \in \Sigma_X$, the procedure* MONITOR($\phi, \mathcal{T}$) *returns*

- $(\top, \varepsilon)$ *if and only if* $\langle\!\langle \phi \rangle\!\rangle(\mathcal{T}) = \top$ *and* $[\![\mathbf{mv}(\phi)]\!](\mathcal{T}) = \varepsilon \geq 0$
- $(\bot, \varepsilon)$ *if and only if* $\langle\!\langle \phi \rangle\!\rangle(\mathcal{T}) = \bot$ *and* $[\![\mathbf{mv}(\phi)]\!](\mathcal{T}) = \varepsilon \leq 0$.

The theoretical complexity of the monitoring algorithms has been studied in the past for both the Linear [23] and the Metric Temporal Logic [10]. Practical algorithms for monitoring using rewriting have been developed by several authors [11, 24]. The new part in Algorithm 2 is the evaluation of the atomic propositions. How easy is to compute the signed distance? When the set X is just $\mathbb{R}$, the set C is an interval and the metric d is the function $d(x, y) = |x - y|$, then the problem reduces to finding the minimum of two values. For example, if $C = [a, b] \subseteq \mathbb{R}$ and $x \in C$, then $\mathbf{Dist}_d(x, C) = \min\{|x - a|, |x - b|\}$. When the set X is $\mathbb{R}^n$, $C \subseteq \mathbb{R}^n$ is a closed and convex set and the metric d is the euclidean distance, i.e. $d(x, y) = ||x - y||_2$, then we can calculate the distance ($\mathbf{dist}_d$) by solving a convex optimization problem. If in addition the set C is a hyperplane $C = \{x \mid a^T x = b\}$

or a halfspace $C = \{x \mid a^T x \leq b\}$, then there exist analytical solutions. For further details see [16].

5 Conclusions and Future Work

The main contribution of this work is the definition of a notion of robust satisfaction of a Linear or Metric Temporal Logic formula which is interpreted over finite timed state sequences that reside in some metric space. We have also presented an algorithmic procedure that can monitor such a timed state sequence and determine an under-approximation of its robustness degree. As mentioned in the introduction, the applications of this framework can extend to several areas. We are currently exploring several new directions such as the extension of the definitions of the robustness degree and the robust MTL semantics so they can handle infinite timed state sequences. Also of interest to us is the addition of a metric on the time bounds as it is advocated in [25] and [26]. Finally, the methodology that we have presented in this paper comprises the basis for the extension of recent results on the safety verification of discrete time systems [13] to a more general verification framework using the metric temporal logic as a specification language [22].

Acknowledgments. The authors would like to thank Oleg Sokolsky, Rajeev Alur, Antoine Girard and Nader Motee for the fruitful discussions and one of the reviewers for the many useful remarks. This work has been partially supported by NSF EHS 0311123, NSF ITR 0324977 and ARO MURI DAAD 19-02-01-0383.

References

1. Clarke, E.M., Grumberg, O., Peled, D.A.: Model Checking. MIT Press, Cambridge, Massachusetts (1999)
2. Alur, R., Courcoubetis, C., Halbwachs, N., Henzinger, T.A., Ho, P.H., Nicollin, X., Olivero, A., Sifakis, J., Yovine, S.: The algorithmic analysis of hybrid systems. Theoretical Computer Science **138** (1995) 3–34
3. Tan, L., Kim, J., Sokolsky, O., Lee, I.: Model-based testing and monitoring for hybrid embedded systems. In: Proceedings of the 2004 IEEE International Conference on Information Reuse and Integration. (2004) 487–492
4. Maler, O., Nickovic, D.: Monitoring temporal properties of continuous signals. In: Proceedings of FORMATS-FTRTFT. Volume 3253 of LNCS. (2004) 152–166
5. Kapinski, J., Krogh, B.H., Maler, O., Stursberg, O.: On systematic simulation of open continuous systems. In: Hybrid Systems: Computation and Control. Volume 2623 of LNCS., Springer (2003) 283–297
6. Esposito, J.M., Kim, J., Kumar, V.: Adaptive RRTs for validating hybrid robotic control systems. In: Proceedings of the International Workshop on the Algorithmic Foundations of Robotics. (2004)
7. Emerson, E.A.: Temporal and modal logic. In van Leeuwen, J., ed.: Handbook of Theoretical Computer Science: Formal Models and Semantics. Volume B., North-Holland Pub. Co./MIT Press (1990) 995–1072

8. Koymans, R.: Specifying real-time properties with metric temporal logic. Real-Time Systems **2** (1990) 255–299
9. de Alfaro, L., Faella, M., Stoelinga, M.: Linear and branching metrics for quantitative transition systems. In: Proceedings of the 31st ICALP. Volume 3142 of LNCS., Springer (2004) 97–109
10. Thati, P., Rosu, G.: Monitoring algorithms for metric temporal logic specifications. In: Runtime Verification. Volume 113 of ENTCS., Elsevier (2005) 145–162
11. Havelund, K., Rosu, G.: Monitoring programs using rewriting. In: Proceedings of the 16th IEEE international conference on Automated software engineering. (2001)
12. Shults, B., Kuipers, B.: Qualitative simulation and temporal logic: proving properties of continuous systems. Technical Report TR AI96-244, Dept. of Computer Sciences, University of Texas at Austin (1996)
13. Girard, A., Pappas, G.J.: Verification using simulation. In: Hybrid Systems: Computation and Control (HSCC). Volume 3927 of LNCS., Springer (2006) 272 – 286
14. Fainekos, G.E., Kress-Gazit, H., Pappas, G.J.: Hybrid controllers for path planning: A temporal logic approach. In: Proceedings of the 44th IEEE Conference on Decision and Control. (2005) 4885 – 4890
15. Lamine, K.B., Kabanza, F.: Reasoning about robot actions: A model checking approach. In: Advances in Plan-Based Control of Robotic Agents. Volume 2466 of LNCS., Springer (2002) 123–139
16. Boyd, S., Vandenberghe, L.: Convex Optimization. Cambridge University Press (2004)
17. Ouaknine, J., Worrell, J.: On the decidability of metric temporal logic. In: 20th IEEE Symposium on Logic in Computer Science (LICS). (2005) 188–197
18. Alur, R., Feder, T., Henzinger, T.A.: The benefits of relaxing punctuality. In: Symposium on Principles of Distributed Computing. (1991) 139–152
19. Alur, R., Henzinger, T.A.: Real-Time Logics: Complexity and Expressiveness. In: Fifth Annual IEEE Symposium on Logic in Computer Science, Washington, D.C., IEEE Computer Society Press (1990) 390–401
20. Fainekos, G.E., Pappas, G.J.: Robustness of temporal logic specifications for finite state sequences in metric spaces. Technical Report MS-CIS-06-05, Dept. of CIS, Univ. of Pennsylvania (2006)
21. Girard, A., Pappas, G.J.: Approximation metrics for discrete and continuous systems. Technical Report MS-CIS-05-10, Dept. of CIS, Univ. of Pennsylvania (2005)
22. Fainekos, G.E., Girard, A., Pappas, G.J.: Temporal logic verification using simulation. In: FORMATS 2006. Volume 4202 of LNCS., Springer (2006) 171–186
23. Markey, N., Schnoebelen, Ph.: Model checking a path (preliminary report). In: Proceedings of the 14th International Conference on Concurrency Theory. Volume 2761 of LNCS. (2003) 251–265
24. Kristoffersen, K.J., Pedersen, C., Andersen, H.R.: Runtime verification of timed LTL using disjunctive normalized equation systems. In: Proceedings of the 3rd Workshop on Run-time Verification. Volume 89 of ENTCS. (2003) 1–16
25. Huang, J., Voeten, J., Geilen, M.: Real-time property preservation in approximations of timed systems. In: Proceedings of the 1st ACM & IEEE International Conference on Formal Methods and Models for Co-Design. (2003) 163–171
26. Henzinger, T.A., Majumdar, R., Prabhu, V.S.: Quantifying similarities between timed systems. In: FORMATS. Volume 3829 of LNCS., Springer (2005) 226–241