Deciding the Satisfiability of MITL Specifications[*]

Marcello M. Bersani*   Matteo Rossi*   Pierluigi San Pietro*,+

* Dipartimento di Elettronica Informazione e Bioingegneria, Politecnico di Milano, Milano, Italy
+CNR IEIIT-MI, Milano, Italy

{marcellomaria.bersani,matteo.rossi,pierluigi.sanpietro}@polimi.it

In this paper we present a satisfiability-preserving reduction from MITL interpreted over finitely-variable continuous behaviors to Constraint LTL over clocks, a variant of CLTL that is decidable, and for which an SMT-based bounded satisfiability checker is available. The result is a new complete and effective decision procedure for MITL. Although decision procedures for MITL already exist, the automata-based techniques they employ appear to be very difficult to realize in practice, and, to the best of our knowledge, no implementation currently exists for them. A prototype tool for MITL based on the encoding presented here has, instead, been implemented and is publicly available.

1 Introduction

Computer systems are inherently discrete-time objects, but their application to control and monitoring of real-time systems often requires to deal with time-continuous external signals and variables, such as position, speed and acceleration or temperature and pressure. Hence, many continuous-time models have been developed for verification and validation of such systems, e.g., Timed Automata [3], or continuous-time temporal logics, such as MITL (Metric Interval Temporal Logic) [4].

In general, the role of temporal logics in verification and validation is two-fold. First, temporal logic allows abstract, concise and convenient expression of required properties of a system. Linear Temporal Logic (LTL) is often used with this goal in the verification of finite-state models, e.g., in model checking [5]. Second, temporal logic allows a descriptive approach to specification and modeling (see, e.g., [19, 14]). A descriptive model is based on axioms, written in some (temporal) logic, defining a system by means of its general properties, rather than by an operational model based on some kind of machine (e.g., a Timed Automaton) behaving in the desired way. In this case, verification typically consists of satisfiability checking of the conjunction of the model and of the (negation of) its desired properties. An example of the latter approach is Bounded Satisfiability Checking (BSC) [20], where Metric Temporal Logic (MTL) specifications on *discrete* time and properties are translated into Boolean logic, in an approach similar to Bounded Model Checking of LTL properties of finite-state machines.

In general, verification of continuous-time temporal logics is not as well sopported as for discrete-time models. Uppaal [6] is the de-facto standard tool for verification of Timed Automata. However, Uppaal does not support continuous-time temporal logics: not only satisfiability checking is not available in Uppaal, but even the formalization of system properties in temporal logic is not allowed, aside from rather simple invariants and reachability properties. Rather, non-trivial properties to be verified on an operational model must be expressed as other Timed Automata, i.e., at a lower level of abstraction. Indeed, there have been a few proposals for verifying continuous-time logics [17], but they do not appear to be actually implementable, and, to the best of our knowledge, in fact they have never been implemented.

This paper proposes a new technique, based on generalizing BSC to MITL, by reducing satisfiability of MITL to satisfiability of *Constraint LTL over clocks* (CLTL-oc), a new decidable variant of CLTL [12].

$$M,t \models p \Leftrightarrow p \in M(t) \qquad p \in AP$$
$$M,t \models \neg\phi \Leftrightarrow M,t \not\models \phi$$
$$M,t \models \phi \wedge \psi \Leftrightarrow M,t \models \phi \text{ and } M,t \models \psi$$
$$M,t \models \phi\mathbf{U}_I\psi \Leftrightarrow \exists t' \in t + I : M,t' \models \psi \text{ and } M,t'' \models \phi \; \forall t'' \in (t,t')$$

Table 1: Semantics of MITL.

In particular, a MITL formula may be encoded into an equisatisfiable CLTL-oc formula, which can then be solved through the same techniques of [7, 9, 8]. The latter approach generalizes BSC to CLTL, generating an encoding suitable for verification with standard Satisfiability Modulo Theories (SMT) solvers such as Z3 [18]. This new technique has been implemented in an open-source prototype tool [1].

Although MITL is known to be decidable over unrestricted behaviors [16], we focus on so-called finitely-variable models, i.e. such that in every bounded time interval there can only be a finite number of changes. This is a very common requirement for continuous-time models, which only rules out pathological behaviors (e.g., Zeno [14]) which do not have much practical interest. To define the encoding, we start by focusing on models in which intervals are closed on the left end and open on the right end. This restriction is later lifted to consider general, finitely-variable, signals.

The paper is organized as follows: Sect. 2 defines MITL and CLTL-oc, Sect. 3 defines a reduction from MITL to CLTL-oc, based on the restriction that intervals are closed to the left and open to the right; Sect. 4 generalizes the translation to intervals of any kind, also discussing the extension to include past operators. Sect. 5 concludes, discussing applications to other logics and presenting a prototype tool.

2   Languages

Let *AP* be a finite set of atomic propositions. The syntax of (well formed) formulae of MITL is defined as follows, with $p \in AP$ and I an interval of the form $\langle a,b \rangle$ or $\langle a,+\infty \rangle$, with $a,b \in \mathbb{N}$ constants, $a < b$:

$$\phi := p \mid \phi \wedge \phi \mid \neg\phi \mid \phi\mathbf{U}_I\phi$$

The semantics of MITL is defined in Table 1 with respect to *signals*. A signal is a function $M : \mathbb{R}_+ \to 2^{AP}$, with $\mathbb{R}_+$ the set of nonnegative reals. A MITL formula ϕ is *satisfiable* if there exists a signal M such that $M,0 \models \phi$ (in this case, M is called a *model* of ϕ). The *globally* $\mathbf{G}_I$ and *eventually* $\mathbf{F}_I$ operators can be defined by the usual abbreviations: $\mathbf{F}_I\phi = \top\mathbf{U}_I\phi$ and $\mathbf{G}_I\phi = \neg\mathbf{F}_I(\neg\phi)$.

Constraint LTL (CLTL [12, 9]) is used in Sect. 3 to solve the satisfiability problem of MITL. CLTL formulae are defined with respect to a finite set V of variables and a *constraint system* $\mathscr{D}$, which is a pair $(D,\mathscr{R})$ with D being a specific domain of interpretation for variables and constants and $\mathscr{R}$ being a family of relations on D, such that the set AP of atomic propositions coincides with set $\mathscr{R}_0$ of 0-ary relations. An *atomic constraint* is a term of the form $R(x_1,\ldots,x_n)$, where R is an n-ary relation of $\mathscr{R}$ on domain D and $x_1,\ldots,x_n$ are variables. A *valuation* is a mapping $v : V \to D$, i.e., an assignment of a value in D to each variable. A constraint is *satisfied* by v, written $v \models_{\mathscr{D}} R(x_1,\ldots,x_n)$, if $(v(x_1),\ldots,v(x_n)) \in R$. Given a variable $x \in V$ over domain D, *temporal terms* are defined by the syntax: $\alpha := c \mid x \mid \mathrm{X}\alpha$, where c is a constant in D and x denotes a variable over D. Operator X is very similar to $\mathbf{X}$, but it only applies to temporal terms, with the meaning that $\mathrm{X}\alpha$ is the *value* of temporal term α in the next time instant.

$$(\pi,\sigma), i \models p \Leftrightarrow p \in \pi(i) \text{ for } p \in AP$$
$$(\pi,\sigma), i \models R(\alpha_1,\ldots,\alpha_n) \Leftrightarrow (\sigma(i+|\alpha_1|,x_{\alpha_1}),\ldots,\sigma(i+|\alpha_n|,x_{\alpha_n})) \in R$$
$$(\pi,\sigma), i \models \neg\phi \Leftrightarrow (\pi,\sigma), i \not\models \phi$$
$$(\pi,\sigma), i \models \phi \wedge \psi \Leftrightarrow (\pi,\sigma), i \models \phi \text{ and } (\pi,\sigma), i \models \psi$$
$$(\pi,\sigma), i \models \mathbf{X}(\phi) \Leftrightarrow (\pi,\sigma), i+1 \models \phi$$
$$(\pi,\sigma), i \models \mathbf{Y}(\phi) \Leftrightarrow (\pi,\sigma), i-1 \models \phi \wedge i > 0$$
$$(\pi,\sigma), i \models \phi\mathbf{U}\psi \Leftrightarrow \exists\, j \geqslant i : (\pi,\sigma), j \models \psi \wedge (\pi,\sigma), n \models \phi \;\forall\, i \leqslant n < j$$
$$(\pi,\sigma), i \models \phi\mathbf{S}\psi \Leftrightarrow \exists\, 0 \leqslant j \leqslant i : (\pi,\sigma), j \models \psi \wedge (\pi,\sigma), n \models \phi \;\forall\, j < n \leqslant i$$

Table 2: Semantics of CLTL.

Well-formed CLTL formulae are defined as follows:

$$\phi := R(\alpha_1,\ldots,\alpha_n) \mid \phi \wedge \phi \mid \neg\phi \mid \mathbf{X}(\phi) \mid \mathbf{Y}(\phi) \mid \phi\mathbf{U}\phi \mid \phi\mathbf{S}\phi$$

where α_i's are temporal terms, $R \in \mathscr{R}$, $\mathbf{X}$, $\mathbf{Y}$, $\mathbf{U}$ and $\mathbf{S}$ are the usual "next", "previous", "until" and "since" operators of LTL, with the same meaning. The dual operators "release" $\mathbf{R}$, and "trigger" $\mathbf{T}$ may be defined as usual, i.e., $\phi\mathbf{R}\psi$ is $\neg(\neg\phi\mathbf{U}\neg\psi)$ and $\phi\mathbf{T}\psi$ is $\neg(\neg\phi\mathbf{S}\neg\psi)$.

The semantics of CLTL formulae is defined with respect to a strict linear order representing time $(\mathbb{N}, <)$. Truth values of propositions in AP, and values of variables belonging to V are defined by a pair (π, σ) where $\sigma : \mathbb{N} \times V \to D$ is a function which defines the value of variables at each position in $\mathbb{N}$ and $\pi : \mathbb{N} \to \wp(AP)$ is a function associating a subset of the set of propositions with each element of $\mathbb{N}$. The value of terms is defined with respect to σ as follows:

$$\sigma(i,\alpha) = \sigma(i+|\alpha|,x_\alpha)$$

where x_α is the variable in V occurring in term α and $|\alpha|$ is the *depth* of a temporal term, namely the total amount of temporal shift needed in evaluating α: $|x| = 0$ when x is a variable, and $|X\alpha| = |\alpha| + 1$. The semantics of a CLTL formula ϕ at instant $i \geqslant 0$ over a linear structure (π, σ) is recursively defined as in Table 2, where $R \in \mathscr{R}\backslash\mathscr{R}_0$. A formula $\phi \in$ CLTL is *satisfiable* if there exists a pair (π, σ) such that $(\pi, \sigma), 0 \models \phi$.

In this paper, we consider a variant of CLTL, where arithmetic variables are evaluated as *clocks* and set $\mathscr{R}$ is $\{<, =\}$. A clock "measures" the time elapsed since the last time the clock was "reset" (i.e., the variable was equal to 0). By definition, in CLTL-oc each $i \in \mathbb{N}$ is associated with a "time delay" $\delta(i)$, where $\delta(i) > 0$ for all i, which corresponds to the "time elapsed" between i and the next state $i+1$. More precisely, for all clocks $x \in V$, $\sigma(i+1,x) = \sigma(i,x) + \delta(i)$, unless it is "reset" (i.e., $\sigma(i+1,x) = 0$).

3 Reduction of MITL to CLTL-over-clocks

This section devises a reduction from MITL to CLTL-oc. The inherent bounded variability of metric operators in MITL allows a translation of a MITL formula ϕ into a CLTL-oc formula with a bounded number of variables, depending on the subformulae of ϕ.

As in [17, 13], it is actually convenient to introduce the operators $\mathbf{U}_{(0,+\infty)}$ and $\mathbf{F}_I$ as primitive, and instead derive the metric until $\mathbf{U}_I$, as shown by the following result.

Lemma 1. *Let M be a signal. Then, for any $t \geqslant 0$,*

$$(1) \quad M,t \models \phi \mathbf{U}_{[a,b\rangle} \psi \Leftrightarrow M,t \models \mathbf{G}_{[0,a)}(\phi \mathbf{U}_{(0,+\infty)} \psi) \wedge \mathbf{F}_{[a,b\rangle} \psi$$

$$(2) \quad M,t \models \phi \mathbf{U}_{(a,b\rangle} \psi \Leftrightarrow M,t \models \mathbf{G}_{[0,a]}(\phi \mathbf{U}_{(0,+\infty)} \psi) \wedge \mathbf{F}_{(a,b\rangle} \psi$$

$$(3) \quad M,t \models \phi \mathbf{U}_{\langle 0,b\rangle} \psi \Leftrightarrow M,t \models \phi \mathbf{U}_{\langle 0,+\infty)} \psi \wedge \mathbf{F}_{\langle 0,b\rangle} \psi$$

When b is $+\infty$, equivalences $(1),(2)$ can be simplified, respectively, in $\phi \mathbf{U}_{[a,+\infty)} \psi \equiv \mathbf{G}_{[0,a)}(\phi \mathbf{U}_{(0,+\infty)} \psi)$ and $\phi \mathbf{U}_{(a,+\infty)} \psi \equiv \mathbf{G}_{[0,a]}(\phi \mathbf{U}_{(0,+\infty)} \psi)$.

The above equivalences make it possible to base the CLTL-oc translation on the $\mathbf{U}_{(0,+\infty)}$ and $\mathbf{F}_I$ operators, instead of $\mathbf{U}_I$, therefore confining metric issues only to the translation of $\mathbf{F}_I$, which is much simpler than the translation of $\mathbf{U}_I$.

Reducing MITL to CLTL-oc requires a way to represent models of MITL formulae, i.e., continuous signals over a finite set of atomic propositions, by means of CLTL-oc models where time is discrete.

Discrete positions in CLTL-oc models represent, for each subformula θ of ϕ, the occurrence of an "event" at that point for the subformula. An "event" is a change of truth value ("become true" or "become false") of θ. Hence, the signal is "stable" (i.e., there is no change) in the interval between two events: a continuous-time signal is hence partitioned by the above events into intervals. Time progress between two discrete points is measured by CLTL variables behaving as clocks: for each subformula θ of ϕ, there are two clocks z_θ^0, z_θ^1 measuring the time elapsed since the last "become true" and "become false" events, respectively (i.e., they are reset when the corresponding event occurs). In case of subformulae of the form $\theta = \mathbf{F}_{\langle a,b\rangle} \phi$, also a finite set of auxiliary clocks is introduced, whose cardinality depends on the values of a, b, namely $d = 2 \left\lceil \frac{b}{b-a} \right\rceil$ auxiliary clocks x_θ^j ($0 \leqslant j \leqslant d-1$). Therefore, a CLTL-oc model embeds, in every (discrete) position both the information defining the truth value of all the subformulae occurring in ϕ and also the time progress between two consecutive events. Then, every position in the CLTL-oc model captures the configuration of one of the intervals in which the MITL signals are partitioned by the events. Therefore, our reduction defines, by means of CLTL-oc formulae, the semantics of every subformula of ϕ.

We start by restricting the set of signals defining models of MITL formulae to signals where intervals are left-closed and right-open (*l.c.r.o.*), e.g.: $\cdots\!\!-\!\!\circ \quad \bullet\!\!-\!\!\cdots$. We will lift this restriction later in the paper. Hence, singularities (i.e., events being true in a single instant) cannot occur and may be ignored. However, the semantics given here does not exclude *a priori* Zeno behaviors [14]: it admits signals corresponding to an infinite sequence of events accumulating to the left of a time instant, i.e., where events do not advance beyond that instant. However, since these signals correspond to behaviors that are of little interest in practice, we restrict the set of models to non-Zeno signals, i.e., to models of CLTL-oc formulae where time diverges: $\sum_{i \in \mathbb{N}} \delta(i) = \infty$, by enforcing a suitable CLTL-oc constraint.

Let M be a signal, ϕ a MITL formula over AP and $sub(\phi)$ the set of all subformulae occurring in ϕ. We write $\uparrow_\theta$ for the occurrence of an event making $\theta \in sub(\phi)$ become true. With abuse of notation we extend $\models$ as follows:

$$M,t \models \uparrow_\theta \Leftrightarrow M,t \models \theta \text{ and } \begin{pmatrix} \exists \varepsilon > 0 \; \forall t' \in (t,t+\varepsilon) \; M,t' \models \theta \text{ and} \\ t > 0 \Rightarrow \exists \varepsilon > 0 \; \forall t' \in (t-\varepsilon,t) \; M,t' \models \neg\theta \end{pmatrix}$$

We define $\downarrow_\theta$ as an abbreviation for $\uparrow_{\neg\theta}$. These definitions impose that signals are defined over an infinite sequence of intervals of the form $[t_1, t_2)$ where $t_2 > t_1$.

Not all temporal operators preserve l.c.r.o. intervals. For example, let $\theta = \mathbf{F}_{\langle a,b\rangle} \phi$ be a MITL formula and let ϕ hold on a l.c.r.o. signal; then, the corresponding signal for θ (i.e., the signal including also the

values for $\uparrow_\theta$), is not l.c.r.o.. In fact, let $t > b$ be the first position such that $M, t \models \uparrow_\phi$. If the signal for θ were l.c.r.o., then it should be $M, t - b \models \uparrow_\theta$, which is impossible because $M, t - b \models \mathbf{F}_{\langle a,b \rangle} \phi \Leftrightarrow \exists t'' \in t - b + \langle a, b \rangle \ M, t'' \models \phi$ and $t'' < t$, but by hypothesis ϕ is false before t. Nevertheless, the next result shows that that Boolean connectives $\neg, \wedge$ and temporal operators $\mathbf{U}_{(0,+\infty)}, \mathbf{F}_{\langle a,b]}, \mathbf{F}_{\langle a,+\infty)}$ and $\mathbf{F}_{\langle 0,b]}$, do indeed preserve l.c.r.o. intervals.

We extend MITL models to any subformulae occurring in MITL formulae by defining a mapping $M_\theta : \mathbb{R}_+ \rightarrow \{\emptyset, \theta\}$ such that:

$$\theta \in M_\theta(t) \Leftrightarrow M, t \models \theta.$$

Lemma 2. *Let M be a l.c.r.o. signal, let ϕ, ψ be two formulae occurring in M and let θ be a formula* $\neg \phi, \phi \wedge \psi, \mathbf{U}_{(0,+\infty)}(\phi, \psi), \mathbf{F}_{\langle a,b]}(\phi), \mathbf{F}_{\langle a,+\infty)}(\phi), \mathbf{F}_{\langle 0,b]}(\phi)$. *Then, M_θ is a l.c.r.o. signal.*

In what follows, $\mathbf{F}_{\langle a,+\infty)}$ is defined as primitive, instead of applying the known equivalence $\mathbf{F}_{[a,+\infty)} \phi \equiv \top \mathbf{U}_{[a,+\infty)} \phi \equiv \mathbf{G}_{[0,a)}(\phi \mathbf{U}_{(0,+\infty)} \psi)$, as formula $\mathbf{G}_{[0,a)} \phi \equiv \neg \mathbf{F}_{[0,a)} \neg \phi$ violates the l.c.r.o. assumption.

We now show how to build a CLTL-oc model (π, σ) of ϕ from a signal M. For each subformula $\theta \in sub(\phi)$ we introduce two *clock variables* z_θ^0, z_θ^1 and one atomic proposition $\overleftarrow{\theta}$. We will ensure that $\overleftarrow{\theta}$ is true at a position whenever θ is true in the interval corresponding to the position. To ease understanding, in the rest we use $\theta = \neg \overleftarrow{\theta}$. We also introduce two abbreviations, $\ulcorner_\theta, \llcorner_\theta$ that play the role of *event markers* (referred to as just "events" when the context is clear); more precisely, they denote, respectively, events $\uparrow_\theta$ and $\downarrow_\theta$, and are defined as follows:

$$\ulcorner_\xi = \neg \mathbf{Y}(\overleftarrow{\xi}) \wedge \overleftarrow{\xi} \qquad \qquad \llcorner_\xi = \neg \mathbf{Y}(\underset{\leftarrow}{\xi}) \wedge \underset{\leftarrow}{\xi}$$

Note that, as $\neg \mathbf{Y}(\bullet)$ is true in the origin, no matter the argument, either $\ulcorner_\theta$ or $\llcorner_\theta$ holds at 0.

For each $\theta = \mathbf{F}_{\langle a,b]} \psi \in sub(\phi)$ we introduce $d = 2 \left\lceil \frac{b}{b-a} \right\rceil$ *auxiliary clocks* $x_\theta^0, \ldots x_\theta^d$. The idea behind the above definitions is that at each occurrence of an event marker ($\ulcorner_\theta$ or $\llcorner_\theta$), exactly one of the clocks z_θ^0, z_θ^1 is equal to 0; the clock, then, measures the time elapsed from the last opposite event. Instead, the auxiliary clocks associated with formulae $\mathbf{F}_{\langle a,b]} \psi$ are used to store the time elapsed since the occurrence of events involving ψ between the current time instant t and $t + b$. In fact, [17] shows that formulae of the form $\mathbf{F}_{\langle a,b \rangle} \psi$ have inherent bounded variability (the result holds for signals with no l.c.r.o. restriction).

Lemma 3 ([17]). *Let $\theta = \mathbf{F}_{\langle a,b \rangle} \psi$, M be a signal and let $0 < t_1 < t_2$ be two instants such that $M, t_1 \models \uparrow_\theta$, $M, t_2 \models \downarrow_\theta$ and $\forall t \in (t', t'') \ M, t \models \theta$. Then, $t_2 - t_1 \geqslant b - a$.*

By Lemma 3, two consecutive events $\uparrow_\theta$ and $\downarrow_\theta$ for formulae $\theta = \mathbf{F}_{\langle a,b \rangle} \psi$ cannot occur at a distance less than $b - a$. However, this does not hold when $\uparrow_\theta$ occurs at $t = 0$ and ψ is true at 0, but it becomes false before b. For instance, let $M, a \models p$ and $M, a + \varepsilon \models \downarrow_p$, where $\varepsilon > 0$ is such that $a + \varepsilon < b$; assume for simplicity that p remains false, i.e., for all $t \in [a + \varepsilon, +\infty), M, t \not\models \psi$. Then, we have that $M, 0 \models \uparrow_\theta$ and $M, \varepsilon \models \downarrow_\theta$. This property will be exploited in Sect. 3.2 to define the translation of the $\mathbf{F}$ operator.

Corollary 1. *Let $\theta = \mathbf{F}_{\langle a,b]} \phi$ be a MITL formula, with $a > 0$, $b \neq \infty$, and let t be an instant of time. Then, in $[t, t + b]$ there are at most $d = 2 \left\lceil \frac{b}{b-a} \right\rceil$ events $\uparrow_\theta, \downarrow_\theta$.*

The result of Corollary 1 can be significantly simplified for formulae of the form $\theta = \mathbf{F}_{\langle 0,b]} \phi$ or of the form $\theta = \mathbf{F}_{\langle a,+\infty)} \phi$. In fact, in the former case, let $t_2 > t_1 \geqslant 0$ be two time instants such that $M, t_1 \models \uparrow_\phi$, $M, t_2 \models \downarrow_\phi$ and $\forall t' \in [t_2, t_2 + b] \ M, t' \not\models \phi$. Then, by definition, we have $M, t_1 - b \models \uparrow_\theta$, $M, t_2 \models \downarrow_\theta$ and $\forall t' \in [t_1 - b, t_2) \models \theta$. Therefore, no event for θ occurs over the interval $[t_1 - b, t_2)$. If $\theta = \mathbf{F}_{\langle a,+\infty)} \phi$, by definition, $M, t \models \theta \Leftrightarrow \exists t' \in \langle t + a, +\infty) \ M, t' \models \phi$; hence, $M, t \models \theta \Rightarrow M, 0 \models \theta$, i.e., $M, 0 \models \uparrow_\theta$. Event

$\uparrow_\theta$ occurs in 0 if, and only if: $\exists t \geqslant a\, M,t \models \uparrow_\phi$ or $\exists t > a\, M,t \models \downarrow_\phi$ or $\exists t < a\, M,t \models \uparrow_\phi \wedge \forall t' > t\, M,t' \models \phi$. Moreover, $M,t \not\models \theta \Rightarrow \forall t' \in \langle t+a, +\infty)\, M,t' \not\models \phi$, i.e., $M,t \models \downarrow_\theta \Leftrightarrow M,t+a \models \downarrow_\theta \wedge \mathbf{G}(\neg\phi)$. By the previous properties, the translation of formulae involving $\mathbf{F}_{\langle 0,b]}$ and $\mathbf{F}_{\langle a,+\infty)}$ is simpler than the case $a > 0$ and $b \neq \infty$, because auxiliary clocks are not needed to represent the formula. For this reason, we provide a direct translation for these subformulae.

Since signals are finitely variable, all the events in M can be enumerated as follows. A position $i \geqslant 0$ uniquely identifies a time instant along M. Let $T \subset \mathbb{R}_+$ be an infinite, but enumerable, set of time instants that includes 0 and every instant when at least one event occurs. Let $I : T \to \mathbb{N}$ be a one-to-one mapping, consistent with the ordering of time, i.e, $I(0) = 0$ and $I(t) < I(t') \Leftrightarrow t < t'$, and such that for all $t_1 < t_2 \in T$ $I(t_2) = I(t_1) + 1 \Leftrightarrow \neg \exists t\, (t_1 < t < t_2 \wedge t \in T)$. By definition, for each subformula θ an event (either $\llcorner_\theta$ or $\lrcorner_\theta$) always occurs at $I(0) = 0$.

Now, given a MITL formula ϕ and a signal M such that $M,0 \models \phi$, we define how to build CLTL-oc interpretations from M. We will prove afterwards that this interpretation is a model for the CLTL-oc formula translating ϕ. We say that a clock v is *reset* at position i when $\sigma(i,v) = 0$.

Let (π,σ) be a CLTL-oc interpretation. If an event for $\theta \in sub(\phi)$ occurs at $t \geqslant 0$, the corresponding event marker ($\llcorner_\theta$ or $\lrcorner_\theta$) labels $\pi(I(t))$ and a reset for one of z_θ^0, z_θ^1 occurs at $I(t)$:

- $\bigvee_{i \in \{0,1\}} \sigma(I(t), z_\theta^i) = 0$ and $(\pi,\sigma), I(t) \models \llcorner_\theta$ if $M,t \models \uparrow_\theta$
- $\bigvee_{i \in \{0,1\}} \sigma(I(t), z_\theta^i) = 0$ and $(\pi,\sigma), I(t) \models \lrcorner_\theta$ if $M,t \models \downarrow_\theta$.
- $\sigma(0, z_\theta^0) = 0$ for all θ.
- $\sigma(0, x_\theta^0) = 0$ for all θ of the form $\mathbf{F}_{\langle a,b \rangle}\psi$.

Note that, by definition, for all time instants $t \in T$ where no events for θ occur, neither $\llcorner_\theta$ nor $\lrcorner_\theta$ hold in $\pi(I(t))$ (i.e., $(\pi,\sigma), I(t) \models \neg\llcorner_\theta \wedge \neg\lrcorner_\theta$).

Now we define how CLTL-oc models represent time progress. Let $t,t' \in T$ be two time instants such that $I(t') = I(t) + 1$. For all clocks z_θ^i that are not reset in $I(t')$ we impose

$$\sigma(I(t'), z_\theta^i) = \sigma(I(t), z_\theta^i) + t' - t.$$

In addition, $\exists i \in \{0,1\}$ s.t. $\sigma(I(t), z_\theta^i) = 0$ if and only if $(\pi,\sigma), I(t) \models \llcorner_\theta$ or $(\pi,\sigma), I(t) \models \lrcorner_\theta$. Clocks z_θ^0, z_θ^1 cannot be reset at the same time, but alternate, and z_θ^0 is reset in the origin. Clocks x_θ^j are dealt with analogously. As mentioned, there exist $d = 2 \left\lceil \frac{b}{b-a} \right\rceil$ clocks x_θ^j for a formula $\mathbf{F}_{\langle a,b \rangle}\psi \in sub(\phi)$. First, for all positions $i \geqslant 0$, $\sigma(i, z_\theta^0) = 0$ or $\sigma(i, z_\theta^1) = 0$ if, and only if, $\bigvee_{j=0}^{d-1} \sigma(i, x_\theta^j) = 0$, i.e, whenever an event for θ occurs, (at least) one auxiliary clock is reset. To avoid simultaneous resets of different clocks, if x_θ^j is reset then no $x_\theta^{j'}$ is reset, for $j' \neq j$. Auxiliary clocks are circularly reset modulo d; i.e., if x_θ^j is reset at position i, then the next reset of x_θ^j, if it exists, occurs in a position $i' > i$ such that all other clocks $x_\theta^{j'}$ ($j' \neq j$) are reset, in order, exactly once in (i,i'). Note that, if a clock x_θ^j is reset at position $i = I(t)$, the next position $i' = I(t')$ when the clock is reset must be such that $t' > t + b$, i.e., given a formula $\theta = \mathbf{F}_{\langle a,b]}$, every clock x_θ^j is reset only once over intervals of length b. The sequence of resets starts with $x_\theta^0 = 0$.

Finally, if ϕ is satisfiable and M is a signal such that $M,0 \models \phi$ i.e., $M,0 \models \uparrow_\phi$, then $(\pi,\sigma), 0 \models \llcorner_\phi$.

Let $r_\phi(M)$ denote the (infinite) set of pairs (π,σ) obtained from M by means of the previous rules for a MITL formula ϕ. The inverse mapping r_ϕ^{-1} is also definable, but not all pairs (π,σ) represent legal signals. Hence, we restrict them to the set of CLTL-oc models that are images of a signal M under r_ϕ, i.e., (π,σ) is such that there exists a signal M such that $(\pi,\sigma) \in r_\phi(M)$. Sect. 3.1 provides a set of CLTL-oc formulae whose models are exactly the set of pairs (π,σ) such that $(\pi,\sigma) \in r_\phi(M)$. For these models the inverse map r^{-1} is well-defined.

3.1 Clocks and Events

The following formulae define how events $\llcorner\theta, \lrcorner\theta$ occur, for $\theta \in sub(\phi)$, and when clocks z_θ^0, z_θ^1 are reset. However, they do not capture the semantics of subformulae θ, which is the object of Sect. 3.2, but only the relations between events $\llcorner\theta$ and $\lrcorner\theta$ and clock resets.

Formula (1) enforces that the occurrence of an event $\llcorner\theta, \lrcorner\theta$ entails the reset of one of z_θ^0, z_θ^1. In addition, Formula $z_\theta^0 = 0$ evaluated in the origin states that clock z_θ^0 is reset in the origin.

$$\llcorner\theta \vee \lrcorner\theta \Leftrightarrow z_\theta^0 = 0 \vee z_\theta^1 = 0 \tag{1}$$

Let $a \in \mathbb{N}$ and value $\bar{a}_k$ be $(a \bmod k)$. The clocks associated with a subformula θ are alternatively reset, as shown on an example in Figure 1. Hence, between any two resets of clock z_θ^0 there must be a reset of clock z_θ^1, and vice-versa:

$$\left(\bigwedge_{i \in \{0,1\}} (z_\theta^i = 0) \right) \Rightarrow \mathbf{X}\left((z_\theta^{\overline{(i+1)_2}} = 0)\mathbf{R}(z_\theta^i \neq 0) \right). \tag{2}$$

For a position $i > 0$ it may happen that neither $\llcorner\theta$ nor $\lrcorner\theta$ occur for any formula (i.e, no events occur). The assumption that intervals are l.c.r.o. entails that intervals have non-null durations, and events $\uparrow_\theta, \downarrow_\theta$ cannot occur at the same time. Define $\mathtt{events}_\phi = \bigwedge_{\theta \in sub(\phi)} (z_\theta^0 = 0) \wedge \mathbf{G}((1) \wedge (2))$.

Lemma 4. *Let θ be a symbol of a MITL formula. For any non-Zeno signal $M : \mathbb{R}_+ \to \{\emptyset, \theta\}$ for θ and for all $(\pi, \sigma) \in r_\theta(M)$, then $(\pi, \sigma), 0 \models \mathtt{events}_\theta$. Conversely, given (π, σ) in which time is divergent and s.t. $(\pi, \sigma), 0 \models \mathtt{events}_\theta$, there is exactly one non-Zeno signal M s.t. $M = r_\theta^{-1}((\pi, \sigma))$.*

Let θ be $\mathbf{F}_{\langle a,b]}\psi$. We introduce $d = 2\left\lceil \frac{b}{b-a} \right\rceil$ clocks x_θ^j, which behave in a similar way as z_θ^0, z_θ^1. Each x_θ^j is needed to store the time elapsed since the occurrence of the last event of θ ($\uparrow_\theta$ or $\downarrow_\theta$). When one of $\uparrow_\theta, \downarrow_\theta$ occurs, then a x_θ^j is reset, i.e., $x_\theta^j = 0$. Each reset event marked by $x_\theta^i = 0$ entails either $\llcorner\theta$ or $\lrcorner\theta$ and all $\uparrow_\theta, \downarrow_\theta$ events are marked by a single reset $x_\theta^i = 0$ (Formula (3)).

$$\left(\llcorner\theta \vee \lrcorner\theta \Leftrightarrow \bigvee_{j=0}^{d-1} x_\theta^j = 0 \right) \wedge \left(\bigwedge_{i=0}^{d-1} \bigwedge_{j=0, i \neq j}^{d-1} \neg(x_\theta^i = 0 \wedge x_\theta^j = 0) \right) \tag{3}$$

The occurrence of resets for clocks x_θ^i is circularly ordered and the sequence of resets starts from the origin by x_θ^0 (see an example in Figure 1). If $x_\theta^i = 0$, then, from the next position, all the other clocks are strictly greater than 0 until the next $x_\theta^{\overline{i+1}_d} = 0$ occurs.

$$\bigwedge_{i=0}^{d-1}\left(x_\theta^i = 0 \Rightarrow \mathbf{X}\left((x_\theta^{\overline{i+1}_d} = 0)\mathbf{R} \bigwedge_{j \in [0,d-1], j \neq i} (x_\theta^{\overline{j+1}_d} > 0) \right) \right) \tag{4}$$

Formula $x_\theta^0 = 0$, evaluated at position 0, sets the first reset of the sequence, constrained by formulae (3)-(4). Moreover, we force all clock values to be strictly ordered in the origin by $x_\theta^0 < x_\theta^{d-1} < \cdots < x_\theta^1$, guaranteeing that resets are correctly associated with events occurring after the origin.

The following lemma (whose proof is similar to the one for Lemma 4) shows that $\mathtt{auxclocks}_\theta$, defined as $(x_\theta^0 = 0) \wedge \mathbf{G}((3) \wedge (4))$ captures map r for $\mathbf{F}_{\langle a,b]}$ formulae. .

Lemma 5. *Let $\theta = \mathbf{F}_{\langle a,b]}\psi$. For any signal $M : \mathbb{R}_+ \to \{\emptyset, \theta\}$ for θ and for all $(\pi, \sigma) \in r_\theta(M)$, it is $(\pi, \sigma), 0 \models \mathtt{auxclocks}_\theta$. Conversely, if $(\pi, \sigma), 0 \models \mathtt{auxclocks}_\theta$, there exists one, and only one, signal M s.t. $M = r_\theta^{-1}((\pi, \sigma))$.*
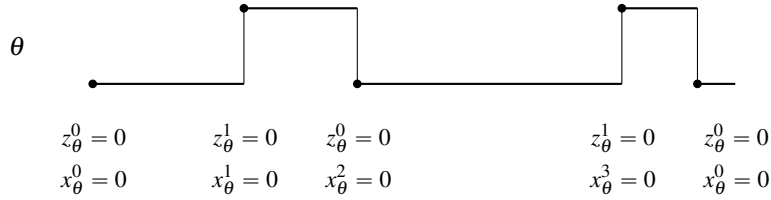
Figure 1: Sequence of circular resets for formula $\theta = \mathbf{F}_{\langle 2,1]}\psi$

3.2   Semantics of MITL Temporal Modalities

We now define a mapping m associating a MITL formula with an equisatisfiable CLTL-oc formula, thus capturing the semantics of MITL in CLTL-oc.

The cases for Boolean connectives and the non-metric $\mathbf{U}$ operator are straightforward. In the following we write O instead of $\neg \mathbf{Y}(\top)$ to represent the first position of CLTL-oc models.

- $\theta = p \in AP$:   it follows from the definition of $\llcorner_p$ and $\lrcorner_p$, representing events $\uparrow_p, \downarrow_p$ over discrete time.

- $\theta = \neg \psi$:   in this case it is $m(\theta) = \overleftarrow{\theta} \Leftrightarrow \overleftarrow{\psi}$.

- $\theta = \gamma \wedge \psi$:   we have: $m(\theta) = \overleftarrow{\theta} \Leftrightarrow \overleftarrow{\gamma} \wedge \overleftarrow{\psi}$.

- $\theta = \gamma \mathbf{U}_{(0,+\infty)}\psi$:   similarly: $m(\theta) = \overleftarrow{\theta} \Leftrightarrow \overleftarrow{\gamma} \wedge \overleftarrow{\gamma} \mathbf{U} \overleftarrow{\psi}$.

- $\theta = \mathbf{F}_{\langle a,b]}\psi$:   When an event $\uparrow_\theta$ occurs, a clock x_θ^j is reset, then event $\uparrow_\psi$ will eventually occur after b time units and it has to occur after $b - a$ instants from the last occurrence of $\downarrow_\psi$ (otherwise $\uparrow_\theta$ has already occurred in the past). The case for $t = 0$ is treated separately: $\uparrow_\theta$ occurs at 0 when there is an interval in which ψ holds that either starts in $[a,b]$ or it spans a. Clock x_θ^0 is used to measure the time elapsing from the origin. In fact, by Corollary 1, x_θ^0, which is reset at 0, can only be reset again after b.

$$
\llcorner_\theta \Leftrightarrow \begin{aligned} &\neg O \wedge \bigvee_{j=0}^{d-1}(x_\theta^j = 0) \wedge \mathbf{X}\left(x_\theta^j > 0 \mathbf{U}\left(\llcorner_\psi \wedge x_\theta^j = b \wedge \bigvee_{i \in \{0,1\}} z_\psi^i > (b-a) \right) \right) \vee \\ &O \wedge (O \vee x_\theta^0 > 0)\mathbf{U}\left(\overleftarrow{\psi} \wedge \left(a \leqslant x_\theta^0 \leqslant b \vee x_\theta^0 < a \wedge \mathbf{X}\left(x_\theta^0 > a \right) \right) \right) \end{aligned} \tag{5}
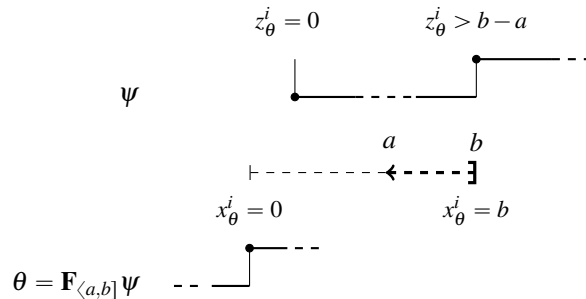$$

Figure 2: Rising edge

Formula (6) defines the condition to make $\llcorner_\theta$ true exactly b instants before an event $\llcorner_\psi$, provided that clock z_ψ^i is greater than $(b-a)$ when $\llcorner_\psi$ occurs (i.e., the last time ψ became false was at least $b-a$

time units before). An illustration of Formulae (5) and (6) is in Figure 2.

$$\ulcorner_\psi \wedge \bigvee_{i\in\{0,1\}} z^i_\psi > (b-a) \Rightarrow \bigvee_{j=0}^{d-1} x^j_\theta = b \tag{6}$$

When an event $\downarrow_\theta$ occurs, a clock x^j_θ is reset, then the event $\downarrow_\psi$ will eventually occur after exactly a time units and the next $\uparrow_\psi$ cannot occur before another $b-a$ instants after that (otherwise $\downarrow_\theta$ cannot occur). In the origin, however, $\downarrow_\theta$ occurs also in the case that $\uparrow_\theta$ does not occur.

$$\lnot_\theta \Leftrightarrow \bigvee_{j=0}^{d-1}(x^j_\theta = 0) \wedge \mathbf{X}\left((x^j_\theta > 0)\mathbf{U}\left(\lnot_\psi \wedge x^j_\theta = a \wedge \ulcorner_\psi\mathbf{R}\lnot\left(\ulcorner_\psi \wedge x^j_\theta \leqslant b\right)\right)\right) \vee (O \wedge \lnot\ulcorner_\theta) \tag{7}$$

Formula (8) is the dual of (6) for a falling edge (Figure 3); it defines a sufficient condition forcing $\lnot_\theta$ when an event $\lnot_\psi$ occurs and $\ulcorner_\psi$ does not happen before $(b-a)$ time units have passed since $\lnot_\psi$.

$$\lnot_\psi \wedge \ulcorner_\psi\mathbf{R}\lnot\left(\ulcorner_\psi \wedge \bigwedge_{i\in\{0,1\}} z^i_\psi \leqslant (b-a)\right) \Rightarrow \bigvee_{j=0}^{d-1} x^j_\theta = a \tag{8}$$
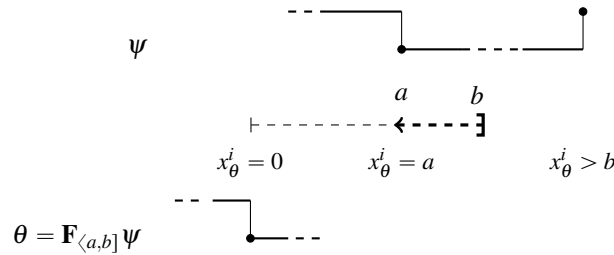

Figure 3: Falling edge

Formula $m(\theta)$ in this case is $(5) \wedge (6) \wedge (7) \wedge (8)$.

As already anticipated, we may study separately the case of formulae $\mathbf{F}_{\langle a,b]}\psi$ where $a = 0$ or $b = +\infty$. The translation in the two cases is simpler than the general one because auxiliary clocks are no longer required to measure the time elapsing between events involving signal for the formula.

- $\theta = \mathbf{F}_{\langle 0,b]}\psi$: the translation for event $\uparrow_\theta$ is analogous to the one of the general case where time elapsing is measured with respect to the clock z^j_θ that is reset when $\ulcorner_\theta$ holds (recall that, by Corollary 1, z^j_θ can be reset again only after the occurrence of $\ulcorner_\psi$). The semantics of $\downarrow_\theta$ in this case is simpler than for Formula (7) because events $\downarrow_\psi$ and $\downarrow_\theta$ always occur simultaneously, provided that the next $\uparrow_\psi$ does not occur within b time instants from $\downarrow_\psi$.

$$\ulcorner_\theta \Leftrightarrow \left(\begin{array}{c} \lnot O \wedge \underleftarrow{\psi} \wedge \left(\bigvee_{j=0}^{1}(z^j_\theta = 0) \wedge \mathbf{X}\left(z^j_\theta > 0\,\mathbf{U}\left(\ulcorner_\psi \wedge z^j_\theta = b \wedge \bigvee_{i\in\{0,1\}} z^i_\psi > b\right)\right)\right) \vee \\ O \wedge (O \vee z^0_\theta > 0)\,\mathbf{U}(\overleftarrow{\psi} \wedge z^0_\theta \leqslant b) \end{array}\right) \tag{9}$$

$$\ulcorner_\psi \wedge \bigvee_{i\in\{0,1\}} z^i_\psi > b \Rightarrow \bigvee_{j\in\{0,1\}} z^j_\theta = b \tag{10}$$

$$\llcorner_\theta \Leftrightarrow \llcorner_\psi \wedge \ulcorner_\psi \mathbf{R} \neg \left(\ulcorner_\psi \wedge \bigwedge_{i \in \{0,1\}} z_\psi^i \leqslant b \right) \tag{11}$$

- $\theta = \mathbf{F}_{\langle a,+\infty \rangle} \psi$: From the semantics of $\mathbf{F}_{a,+\infty}(\psi)$ it is easy to see that event $\uparrow_\theta$ may only occur at 0, if ψ eventually holds in the future after a instants from the origin. Similarly, event $\downarrow_\theta$ may only occur once, but not necessarily in the origin; more precisely, it holds at 0 if and only if $\uparrow_\theta$ does not hold at 0, while for every instant $t > 0$ it occurs when event $\downarrow_\psi$ occurs in $t + a$ and ψ is always false afterwards. As a consequence, z_θ^1 is reset at most once, if $\llcorner_\theta$ occurs in an instant other than the origin.

$$\ulcorner_\theta \Leftrightarrow O \wedge (O \vee z_\theta^0 > 0) \mathbf{U} \left(\overleftarrow{\psi} \wedge \left(a \leqslant z_\theta^0 \vee z_\theta^0 < a \wedge \mathbf{X}\left(z_\theta^0 > a\right) \right) \right) \tag{12}$$

$$\llcorner_\theta \Leftrightarrow z_\theta^1 = 0 \wedge \mathbf{X}\left(z_\theta^1 > 0 \mathbf{U}\left(\llcorner_\psi \wedge z_\theta^1 = a \wedge \mathbf{G}\left(\neg \ulcorner_\psi\right)\right)\right) \vee (O \wedge \neg \ulcorner_\theta) \tag{13}$$

$$\llcorner_\psi \wedge \mathbf{G}\left(\neg \ulcorner_\psi\right) \Rightarrow z_\theta^1 = a \tag{14}$$

3.3 Correctness

Let F be a set of formulae. We extend map r to $sub(\phi)$, written $r_{sub(\phi)}(M)$, to represent the set of CLTL-oc models where atomic propositions are symbols associated with each subformula in ϕ and variables includes all clocks z_θ^0, z_θ^1 and the auxiliary clocks for the case $\mathbf{F}_{\langle a,b \rangle}$.

Lemma 6. *Let M be a signal, and ϕ a MITL formula. For any $(\pi, \sigma) \in r_{sub(\phi)}(M)$ it is:*

$$(\pi, \sigma), 0 \models \bigwedge_{\theta \in sub(\phi)} \mathbf{G}(m(\theta)) \wedge \mathtt{events}_\theta \wedge \bigwedge_{\substack{\theta \in sub(\phi) \\ \theta = \mathbf{F}_{\langle a,b \rangle}}} \mathtt{auxclocks}_\theta$$

and for all $k \in \mathbb{N}, \theta \in sub(\phi)$ it is $(\pi, \sigma), k \models m(\theta)$.

Lemma 7. *Let M be a signal and let ϕ be a MITL formula. If*

$$(\pi, \sigma), 0 \models \bigwedge_{\theta \in sub(\phi)} \mathbf{G}(m(\theta)) \wedge \mathtt{events}_\theta \wedge \bigwedge_{\substack{\theta \in sub(\phi) \\ \theta = \mathbf{F}_{\langle a,b \rangle}}} \mathtt{auxclocks}_\theta$$

and $M = r_{sub(\phi)}^{-1}((\pi, \sigma))$, then for all $t \in T$ it is $(\pi, \sigma), I(t) \models \ulcorner_\phi$ iff $M, t \models \uparrow_\phi$ (similarly for $\llcorner_\phi$).

The main result, the equisatisfiability of MITL and of its CLTL-oc translation, follows.

Theorem 1. *A MITL formula ϕ is satisfiable if, and only if the following formula is satisfiable:*

$$\ulcorner_\phi \wedge \bigwedge_{\theta \in sub(\phi)} \mathbf{G}(m(\theta)) \wedge \mathtt{events}_\phi \wedge \bigwedge_{\substack{\theta \in sub(\phi) \\ \theta = \mathbf{F}_{\langle a,b \rangle}}} \mathtt{auxclocks}_\theta. \tag{15}$$

3.4   Complexity

The reduction of MITL to CLTL-oc of Sect. 3.2 induces an EXPSPACE decision procedure for the satisfiability of MITL (the problem is actually EXPSPACE-complete). In fact, consider a MITL formula φ, and its CLTL-oc translation (15) obtained following the reduction of Sect. 3.2. In Formula (15) we introduce two clocks for each subformula of φ, unless the subformula is of the form $\mathbf{F}_{\langle a,b]}\psi$, in which case we introduce at most b clocks, since $a,b \in \mathbb{N}$. Then, the size of (15) is $O(|\varphi|K)$, where K is the maximum constant appearing in φ. It can be shown that satisfiability for a CLTL-oc formula ϕ_{CLTL} is PSPACE in the number of subformulae of ϕ_{CLTL} (which is $O(|\varphi|K)$ for Formula (15)) and in the size of the string encoding the maximum constant occurring in it (K for Formula (15)). Hence, the decision procedure induced by our encoding is in EXPSPACE when using a binary encoding of K. As remarked in [4], if the MITL formula φ does not contain subformulae of type $\mathbf{F}_{\langle a,b]}\psi$ (with $a > 0$ and $b \neq \infty$), the reduction of Sect. 3.2 only introduces one clock variable for each subformula. As a consequence, the size of Formula (15) is $O(|\varphi|)$ and the algorithm is in PSPACE.

4   Generalized translation

Our translation from MITL to CLTL-oc can be extended to represent general signals where no assumption is made on their shape, other than their finite variability, i.e., the l.c.r.o. assumption of Sect. 3 can be relaxed. In this more general case, the truth of a formula ϕ can change in a *singular* manner, that is, there can be instants where the value of ϕ is different than in a neighborhood thereof.

   More precisely, we say that in a time instant t of a signal M formula ϕ has an "up-singularity" s_ϕ^u if it holds in t, but not before and after it; more precisely, we say that $M,t \models s_\phi^u$ if and only if $M,t \models \phi$ and $\exists \varepsilon > 0$ s.t. $\forall t' \neq t \in (t - \varepsilon, t + \varepsilon)$ it is $M,t' \not\models \phi$. We say that ϕ has a "down-singularity" s_ϕ^d when $\neg\phi$ has an up-singularity (i.e., ϕ does not hold in t, but it does before and after it). Note that, by their definition, singularities (either up or down), cannot occur in $t = 0$.

   To represent general signals in CLTL-oc we "split" the representation of the value of subformulae θ in intervals $[t,t')$ in two parts: $\uparrow_\theta$ captures the value of θ in t, whereas $\overleftarrow{\theta}$ corresponds to its value in (t,t'). With the new predicates, we can restrict represented signals to only include l.c.r.o. intervals by imposing the constraint $\quad \uparrow_\theta \Leftrightarrow \overleftarrow{\theta} \quad$ for all θ. In addition, $\overline{\theta}$ and $\underline{\theta}$ become: $\quad \overline{\theta} = \uparrow_\theta \wedge \overleftarrow{\theta} \qquad \underline{\theta} = \neg\,\uparrow_\theta \wedge \neg\,\overleftarrow{\theta}$. Then, the encoding of Sect. 3 can be used also with the new atomic predicates, provided constraint $\uparrow_\theta \Leftrightarrow \overleftarrow{\theta} \quad$ is added for all subformulae. If, instead, general signals are to be allowed, the encoding must be extended to include also the cases in which the values of (sub)formulae change in singular manners.

   To this end, we slightly modify the definition of $\llcorner_\xi$ as $\neg\mathbf{Y}(\overleftarrow{\xi}) \wedge \overleftarrow{\xi}$ and $\llcorner_\xi$ as $\neg\mathbf{Y}(\neg\overleftarrow{\xi}) \wedge \neg\overleftarrow{\xi}$ and we introduce the following abbreviations, which capture, respectively, up- and down-singularities (note that neither $\text{\textsl{Ɯ}}_\xi$, nor $\text{\textsf{T}}_\xi$ hold at 0, as $\mathbf{Y}(\bullet)$ is false there):

$$\text{\textsl{Ɯ}}_\xi = \mathbf{Y}(\neg\overleftarrow{\xi}) \wedge \uparrow_\xi \wedge \neg\overleftarrow{\xi} \qquad \text{\textsf{TΓ}}_\xi = \mathbf{Y}(\overleftarrow{\xi}) \wedge \neg\,\uparrow_\xi \wedge \overleftarrow{\xi}$$

We also define the following: $\qquad \overset{\xi}{\rightharpoonup} = \ulcorner_\xi \vee \text{\textsl{Ɯ}}_\xi \vee (O \wedge \uparrow_\xi) \qquad \overset{\xi}{\rightharpoondown} = \lrcorner_\xi \vee \text{\textsl{Ɯ}}_\xi.$

More precisely, $\overset{\xi}{\rightharpoonup}$ corresponds to a situation where ξ does not hold the interval before the current one (if such interval exists), and it is true sometimes in the current one (either in its first instant, in which case ξ

can have a up-singularity, or in the rest of the interval). Dually, $\overset{\xi}{\hookleftarrow}$ holds if ξ is true in the first instant of the current interval, or in the interval before it, and from that moment on it is false.

When general signals are allowed, there is no need to restrict the temporal operators only to $\mathbf{F}_{\langle a,b]}(\psi)$. For simplicity, we focus on the encoding of case $\theta = \mathbf{F}_{(a,b)}(\psi)$, all other cases being similar.

- $\theta = \mathbf{F}_{(a,b)}\psi$: We have the following result.

Lemma 8. *If $\theta = \mathbf{F}_{(a,b)}\psi$ is a MITL formula and $M,t \models \theta$ then s $\exists \varepsilon \in \mathbb{R}_{>0}$ such that, for all $t' \in [t,t+\varepsilon]$ it is $M,t' \models \theta$ and, when $t > 0$, there is also $\varepsilon \in \mathbb{R}_{>0}$ such that $\varepsilon < t$ and for all $t' \in [t-\varepsilon,t]$ it is $M,t' \models \theta$.*

Because of Lemma 8, an up-singularity $\rfloor\llcorner_\theta$ can never occur for $\theta = \mathbf{F}_{(a,b)}\psi$. In addition, if θ holds at the beginning of an interval (i.e., $\upharpoonleft_\theta$ holds), then it must hold also in the rest of the interval and, if $t > 0$, it must also hold in the interval before. Then, the following constraint holds in every instant:

$$\upharpoonleft_\theta \Rightarrow \overset{\leftarrow}{\theta} \wedge (\mathbf{Y}(\overset{\leftarrow}{\theta}) \vee O) \tag{16}$$

Formula (17) is similar to (5), but it specifies that, when θ becomes true outside of the origin, it must do so in a left-open manner (i.e., $\upharpoonleft_\theta$ does not hold with $\rfloor\urcorner_\theta$); also, there is one additional condition that makes θ become true in 0 when ψ becomes true exactly at b, in which case θ does not hold in 0.

$$\rfloor\urcorner_\theta \Leftrightarrow \begin{array}{l} \neg O \wedge \neg \upharpoonleft_\theta \wedge \bigvee_{j=0}^{d-1}(x_\theta^j = 0) \wedge \mathbf{X}\left(x_\theta^j > 0 \mathbf{U}\left(\overset{\psi}{\urcorner} \wedge x_\theta^j = b \wedge \bigvee_{i=0}^{1} z_\psi^i > (b-a)\right)\right) \quad \vee \\[2em] O \wedge \neg \upharpoonleft_\theta \wedge \mathbf{X}\left(x_\theta^0 > 0 \mathbf{U}\left(\overset{\psi}{\urcorner} \wedge x_\theta^0 = b \wedge \bigvee_{i=0}^{1} z_\psi^i \geq (b-a)\right)\right) \quad \vee \\[2em] O \wedge \upharpoonleft_\theta \wedge (O \vee x_\theta^0 > 0)\mathbf{U}\left((\upharpoonleft_\psi \vee \overset{\leftarrow}{\psi}) \wedge a < x_\theta^0 < b \quad \vee \quad \overset{\leftarrow}{\psi} \wedge x_\theta^0 < a \wedge \mathbf{X}(x_\theta^0 > a)\right) \end{array} \tag{17}$$

Formulae (18), (19) and (20) generalize, respectively, (6), (7) and (8) to include also the case in which ψ changes its value in a singular manner (i.e., with $\rfloor\llcorner_\psi$ instead of $\rfloor\urcorner_\psi$ or $\urcorner\llcorner_\psi$).

$$\overset{\psi}{\urcorner} \wedge \bigvee_{i\in\{0,1\}} z_\psi^i \geq (b-a) \Rightarrow \bigvee_{j=0}^{d-1} x_\theta^j = b \tag{18}$$

$$\urcorner\llcorner_\theta \Leftrightarrow \bigvee_{j=0}^{d-1}(x_\theta^j = 0) \wedge \mathbf{X}\left((x_\theta^j > 0)\mathbf{U}\left(\overset{\psi}{\hookleftarrow} \wedge x_\theta^j = a \wedge \mathbf{X}\left(\overset{\psi}{\urcorner}\mathbf{R}\neg\left(\overset{\psi}{\urcorner} \wedge x_\theta^j \leq b\right)\right)\right)\right) \quad \vee \quad (O \wedge \neg\rfloor\urcorner_\theta) \tag{19}$$

$$\overset{\psi}{\hookleftarrow} \wedge \mathbf{X}\left(\overset{\psi}{\urcorner}\mathbf{R}\neg\left(\overset{\psi}{\urcorner} \wedge \bigwedge_{i=0}^{1} z_\psi^i \leq (b-a)\right)\right) \Rightarrow \bigvee_{j=0}^{d-1} x_\theta^j = a \tag{20}$$

Finally, we need to consider an additional shape in which θ can change value. More precisely, there is also the case in which θ becomes false with a down-singularity $\overline{\urcorner\llcorner}_\theta$. This occurs in an instant t (which must be > 0, as singularities cannot occur in the origin by definition) such that ψ becomes false at $t + a$, but it becomes true again at $t + b$ (and it stays false in interval $(t + a, t + b)$). This condition is captured by Formula (21), which is similar to Formula (19), except that it specifies that when ψ becomes true again, the clock x_θ^j that is reset when ϕ has the singularity has value b.

$$\overline{\urcorner\llcorner}_\theta \Leftrightarrow \neg O \wedge \bigvee_{j=0}^{d-1}(x_\theta^j = 0) \wedge \mathbf{X}\left((x_\theta^j > 0)\mathbf{U}\left(\overset{\psi}{\hookleftarrow} \wedge x_\theta^j = a \wedge \mathbf{X}\left(\overset{\psi}{\cancel{\urcorner}}\mathbf{U}\left(\overset{\psi}{\urcorner} \wedge x_\theta^j = b\right)\right)\right)\right) \tag{21}$$

Then, $m(\theta)$ is (16) $\wedge$ (17) $\wedge$ (18) $\wedge$ (19) $\wedge$ (20) $\wedge$ (21).

To allow for signals of general shape, the encoding for subformulae of the form $\gamma \mathbf{U}_{(0,+\infty)} \psi$ must also be revisited. As this is rather straightforward, we skip the details for reasons of brevity. Instead, we point out that it is possible to define a CLTL-oc encoding also for MITL *past* operators $\mathbf{S}$ and $\mathbf{P}_{\langle a,b \rangle}$. It is known that past operators increase the expressiveness of MITL [11], but do not impact on decidability. Hence, a decision procedure that also includes the possibility to handle past operators is more powerful than one dealing with the future-only fragment. To conclude this section, we show the encoding $m(\theta)$ for the $\mathbf{S}$ operator (whose semantics is symmetric to the one of $\mathbf{U}$ shown in Table 1). The case for operator $\mathbf{P}_{\langle a,b \rangle}$ is omitted for brevity.

- $\theta = \gamma \mathbf{S}_{(0,+\infty)} \psi$:   In this case it can be shown that, if M is a finitely variable signal and θ holds in an instant t, then it must also hold in $(t - \varepsilon, t)$, for some $\varepsilon > 0$, and vice-versa. Then, in $t = 0$ θ is false, and there $\mathbf{S}$ formulae cannot have singularity points. In addition, when a $\mathbf{S}$ formula changes its value after the origin, it must do so in a left-open manner (i.e., the value at the changing point is the same as the one before the changing point). Then, we have

$$m(\theta) = (\mathord{\uparrow}_\theta \Leftrightarrow \mathbf{Y}(\overset{\leftarrow}{\theta})) \wedge (\overset{\leftarrow}{\theta} \Leftrightarrow \overset{\leftarrow}{\gamma} \mathbf{S}((\mathord{\uparrow}_\psi \vee \overset{\leftarrow}{\psi}) \wedge \overset{\leftarrow}{\gamma})). \tag{22}$$

5   Conclusions

This paper investigates a bounded approach to satisfiability checking of the continuous-time temporal logic MITL. We showed an encoding of MITL into a decidable logic (CLTL-oc), which allows, both in principle and in practice, the use of SMT solvers to check satisfiability of MITL.

A decision procedure for CLTL-oc [10] is implemented in a plugin, called `ae`^{2}`zot`, of our Zot toolkit [2], whereas the reduction outlined in Sect. 3 and 4 is implemented in the `qtlsolver` tool, available from [1]. The tool translates MITL (or the expressively equivalent QTL logic [16]) into CLTL-oc, which can be checked for satisfiability by `ae`^{2}`zot`. The resulting toolkit has a 3-layered structure, where CLTL-oc is the intermediate layer between SMT-solvers and various temporal formalisms that can be reduced to CLTL-oc. This not only supports (bounded) satisfiability verification of different languages, but it also allows the expression of different degrees of abstraction. For instance, MITL abstracts away the notion of clocks, inherently encompassed within temporal modalities, which are instead explicit in CLTL-oc and actually available to a user, e.g., to express or verify properties where clocks are convenient. In fact, preliminary experimental results point out that the time required to solve CLTL-oc may be significantly smaller than the one needed for more abstract languages, such as MITL. This is caused by the "effort" required to capture the semantics of temporal modalities, which, on the other hand, allow for more concise and manageable high-level specifications. This layered structure also allows the resolution of a formula to be compliant with constraints imposed at lower layers, for instance by adding at the CLTL-oc layer some extra formula limiting the set of valid models (e.g., by discarding certain edges of some events or by adding particular timing requirements). Also the third layer (the SMT solver) may be used to add further constraints, e.g., to force the occurrence of a proposition or of a certain clock value at a specific discrete position of the finite model.

The current implementation of `qtlsolver` supports the MITL-to-CLTL-oc translation, both with or without the l.c.r.o. restriction. In fact, the following encodings are currently available:

MITL providing a direct definition of MITL operators, assuming l.c.r.o. intervals;

QTL providing the definition of generalized QTL operators (e.g., $\mathbf{F}_{(0,b)}$, $\mathbf{P}_{(0,b)}$) with unrestricted signals (other than they be finitely variable), and MITL operators through abbreviations.

We used the above two encodings to carry out some experiments (available from the `qtlsolver` website [1], or described in [10]). Let us illustrate one of them. MITL Formula (23) specifies that predicate p occurs in isolated points with a period of 100 (i.e., it occurs exactly at 0, 100, 200, etc.).

$$\mathbf{G}_{[0,\infty)} \left(\left(\mathbf{G}_{(0,100)}(\neg p) \Rightarrow \mathbf{G}_{(100,200)}(\neg p) \right) \wedge \left(p \Rightarrow \mathbf{F}_{(0,200)}(p) \right) \right) \wedge p \wedge \mathbf{G}_{(0,100)}(\neg p) \qquad (23)$$

`qtlsolver` was able to find a model for Formula (23) in around 10 seconds, using a bound of 10.[1] Note that, even if the constants appearing in Formula (23) are in the order of the hundreds, events in the corresponding models occur only sparsely, hence a bound of 10 is enough for `qtlsolver` to satisfy (23). If we add to the specification Formula (24), which states that q must hold within 1 time unit in the past or in the future of each p, the solver finds a model (again, with bound 10) in about 40 seconds.

$$\mathbf{G}_{(0,\infty)} \left(p \Rightarrow \mathbf{F}_{(0,1)}(q) \vee \mathbf{P}_{(0,1)}(q) \right) \qquad (24)$$

Formula (24) does not impose that q be false in between occurrences of p. A more restricted behavior is obtained by adding also constraint (25), which imposes that q occurs only in isolated instants, and that there must be at least 100 time units between consecutive occurrences of q.

$$\mathbf{G}_{(0,\infty)} \left(q \Rightarrow \mathbf{G}_{(0,100)}(\neg q) \right) \qquad (25)$$

`qtlsolver` was able to find a model (with bound 20, in this case) for formula (23) $\wedge$ (24) $\wedge$ (25) in around 10 minutes. As mentioned above, one can add constraints at different levels of abstraction. For example, we can add SMT constraints imposing that the *values* of the clocks (instead of the clock regions) associated with propositions p and q be periodic; this allows us to check that formula (23) $\wedge$ (24) $\wedge$ (25) admits periodic models (`qtlsolver` takes around 15 minutes to produce one with bound 20). Finally, if in Formula (25) we replace $\mathbf{G}_{(0,100)}$ with $\mathbf{G}_{(0,100]}$, the behavior becomes strictly aperiodic. In this case the solver takes around 80 minutes to find a model with bound 30, and in excess of 12 hours to show that, with that bound, no model exists in which p and q are periodic (i.e., that the specification, with the added constraint that the values of the clocks associated with p and q be periodic, is unsatisfiable).

While the results presented above are promising, further research will focus on optimizing the implementation of the solver and on extending the encoding to deal with richer constraints.

The techniques presented in this paper for MITL can be tailored also to other logics. We consider an example here. A syntactic fragment of MITL was proposed in [15], namely $\text{MTL}_{0,\infty}$, where temporal modalities are restricted only to intervals of the form $\langle 0,b \rangle$ or $\langle a,\infty \rangle$ (e.g., the MITL formula $\mathbf{F}_{(2,3)}\phi$ is not acceptable). $\text{MTL}_{0,\infty}$ is complete in the sense that every MITL formula can be transformed into an equisatisfiable $\text{MTL}_{0,\infty}$ formula. However, the transformation may lead to an exponential blow-up, since satisfiability is EXPSPACE-complete for MITL and PSPACE-complete for $\text{MTL}_{0,\infty}$. In [15], $\text{MTL}_{0,\infty}$ was shown to be equivalent to a new temporal logic, called Event-Clock Logic (ECL), which is also in PSPACE. Although our work only concerns MITL (and actually $\text{MTL}_{0,\infty}$, which is considered by our translation provided that operator $\mathbf{F}_{\langle a,b]}$ is not primitive for the language), our results can directly be applied for solving the satisfiability of ($\text{MTL}_{0,\infty}$ and) ECL as well, by means of the above equivalence of the languages. However, an explicit encoding of ECL into CLTL-oc may be devised, since only a finite number of explicit clocks are enough to capture ECL semantics; this may allow solving satisfiability of both logics ($\text{MTL}_{0,\infty}$ and ECL) in PSPACE.

---

[1] All tests have been carried out on a desktop computer with a 2.8GHz AMD PhenomTMII processor and 8MB RAM; the solver was Microsoft Z3 3.2. The encoding used was the one for QTL, with unrestricted signals.

References

[1] *qtlsolver*. available from `qtlsolver.googlecode.com`.

[2] *Zot: a Bounded Satisfiability Checker*. available from `zot.googlecode.com`.

[3] Rajeev Alur & David L. Dill (1994): *A theory of timed automata*. Theor. Comp. Sci. 126(2), pp. 183–235. Available at `http://dx.doi.org/10.1016/0304-3975(94)90010-8`.

[4] Rajeev Alur, Tomás Feder & Thomas A. Henzinger (1996): *The Benefits of Relaxing Punctuality*. Journal of the ACM 43(1), pp. 116–146. Available at `http://doi.acm.org/10.1145/112600.112613`.

[5] Christel Baier & Joost-Pieter Katoen (2008): *Principles of Model Checking*. MIT Press.

[6] Johan Bengtsson & Wang Yi (2004): *Timed Automata: Semantics, Algorithms and Tools*. In: *Lect. on Concurrency and Petri Nets*, LNCS 3098, Springer, pp. 87–124. Available at `http://dx.doi.org/10.1007/978-3-540-27755-2_3`.

[7] Marcello M. Bersani, Achille Frigeri, Angelo Morzenti, Matteo Pradella, Matteo Rossi & Pierluigi San Pietro (2010): *Bounded Reachability for Temporal Logic over Constraint Systems*. In: *TIME*, IEEE Computer Society, pp. 43–50. Available at `http://dx.doi.org/10.1109/TIME.2010.21`.

[8] Marcello M. Bersani, Achille Frigeri, Angelo Morzenti, Matteo Pradella, Matteo Rossi & Pierluigi San Pietro (2012): *CLTL Satisfiability Checking without Automata*. arXiv:1205.0946v1.

[9] Marcello M. Bersani, Achille Frigeri, Matteo Rossi & Pierluigi San Pietro (2011): *Completeness of the Bounded Satisfiability Problem for Constraint LTL*. In: *Reachability Problems*, LNCS 6945, pp. 58–71. Available at `http://dx.doi.org/10.1007/978-3-642-24288-5_7`.

[10] Marcello M. Bersani, Matteo Rossi & Pierluigi San Pietro (2013): *A Tool for Deciding the Satisfiability of Continuous-time Metric Temporal Logic*. In: *Proceedings of the International Symposium on Temporal Representation and Reasoning (TIME)*. To appear.

[11] Patricia Bouyer, Fabrice Chevalier & Nicolas Markey (2010): *On the expressiveness of TPTL and MTL*. Information and Computation 208(2), pp. 97 – 116. Available at `http://dx.doi.org/10.1016/j.ic.2009.10.004`.

[12] Stéphane Demri & Deepak D'Souza (2007): *An automata-theoretic approach to constraint LTL*. Information and Computation 205(3), pp. 380–415. Available at `http://dx.doi.org/10.1016/j.ic.2006.09.006`.

[13] Deepak D'Souza & Nicolas Tabareau (2004): *On Timed Automata with Input-Determined Guards*. In: *Proc. of FORMATS/FTRTFT*, LNCS 3253, Springer, pp. 68–83. Available at `http://dx.doi.org/10.1007/978-3-540-30206-3_7`.

[14] Carlo A. Furia, Dino Mandrioli, Angelo Morzenti & Matteo Rossi (2012): *Modeling Time in Computing*. EATCS Monographs in Theoretical Computer Science, Springer. Available at `http://dx.doi.org/10.1007/978-3-642-32332-4`.

[15] Thomas A. Henzinger, Jean F. Raskin & Pierre Y. Schobbens (1998): *The Regular Real-Time Languages*. In: *Proc. of ICALP'98*, LNCS 1343, pp. 580–591. Available at `http://dx.doi.org/10.1007/BFb0055086`.

[16] Yoram Hirshfeld & Alexander Moshe Rabinovich (2004): *Logics for Real Time: Decidability and Complexity*. Fundamenta Informaticae 62(1), pp. 1–28.

[17] Oded Maler, Dejan Nickovic & Amir Pnueli (2006): *From MITL to Timed Automata*. In: *Proc. of FORMATS*, LNCS 4202, pp. 274–289. Available at `http://dx.doi.org/10.1007/11867340_20`.

[18] Microsoft Research (2009): *Z3: An Efficient SMT Solver*. Available at: http://research.microsoft.com/en-us/um/redmond/projects/z3/.

[19] Angelo Morzenti & Pierluigi San Pietro (1994): *Object-Oriented Logical Specification of Time-Critical Systems*. ACM Transactions on Software Engineering and Methodology (TOSEM) 3(1), pp. 56–98. Available at `http://doi.acm.org/10.1145/174634.174636`.

[20] Matteo Pradella, Angelo Morzenti & Pierluigi San Pietro (2013): *Bounded Satisfiability Checking of Metric Temporal Logic Specifications*. ACM Trans. on Soft. Eng. and Meth. (TOSEM). To appear.