



Proceedings of the
Automated Verification of Critical Systems
(AVoCS 2013)

On the Satisfiability of Metric Temporal Logics over the Reals

Marcello M. Bersani, Matteo Rossi and Pierluigi San Pietro

15 pages

On the Satisfiability of Metric Temporal Logics over the Reals

Marcello M. Bersani¹, Matteo Rossi¹ and Pierluigi San Pietro^{1,2}

¹ [marcellomaria.bersani,matteo.rossi,pierluigi.sanpietro]@polimi.it

Dipartimento di Elettronica Informazione e Bioingegneria, Politecnico di Milano

² CNR IEIIT-MI

Abstract: We show that there is a satisfiability-preserving translation of QTL formulae interpreted over finitely variable behaviors into formulae of the CLTL-over-clocks logic. The satisfiability of CLTL-over-clocks can be determined through a suitable encoding into the input logics of SMT solvers, so it constitutes an effective decision procedure for QTL. Although decision procedures for determining satisfiability of QTL (and for the expressively equivalent logics MITL and QMLO) 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 QTL based on the encoding presented here has, instead, been implemented and is publicly available.

Keywords: Metric Temporal Logic, Satisfiability Modulo Theories, Continuous-time systems, Formal Verification

1 Introduction

The need for continuous-time models arises naturally and often when describing the dynamics of physical quantities, such as position, speed and acceleration of a moving body, or such as temperature and pressure of a fluid. When developing computer systems that monitor and control such quantities, then, the classic discrete-time models used in the computer science domain are no longer enough. Many notations [FMMR12] have been developed to address these shortcomings; the most successful ones, those that are the most used in practice and with the most developed tools, are based on operational mechanisms, e.g., Timed Automata [AD94].

Descriptive notations, e.g., temporal logics, however, provide many benefits, such as allowing for an abstract, concise and convenient expression of the required properties of a system. This is mostly exploited in the verification of finite-state models, e.g., through model checking [BK08]. Temporal logics, however, also allow designers to pursue a descriptive approach to the specification and modeling of reactive systems (e.g., [MP94, FMMR12]), where the system is defined by its general properties, rather than by a machine behavior (e.g., a Timed Automaton). In this case, verification typically consists of satisfiability checking of the conjunction of the model and of the (negation of) its desired properties.

In general, tool support for verification of continuous-time temporal logics is not as well-developed as for discrete-time models, especially when the logic is endowed with metric operators. Decision procedures for determining the satisfiability of continuous-time metric temporal logic mostly rely on timed automata-based techniques [AFH96, MNP06], but they appear to be very difficult to realize in practice, and, to the best of our knowledge, no implementation exists

for them. An alternative proof to the one in [AFH96] for the satisfiability of Metric Interval Temporal Logic (MITL) formulae is provided in [SRH02]. Though the aim of the paper was that of proving the soundness and completeness of the axiomatization for the Event-Clock logic (therein proved to be equivalent to MITL), they devise an ad-hoc procedure for building an automaton corresponding to a formula, motivating it since the known one for MITL [AFH96] can not be used directly for their purposes.

We study the satisfiability of the Quantitative Temporal Logic (QTL) [HR99, HR05], using a purely logic-based approach. QTL is an interesting logic: it is known to be decidable over the real line, and its satisfiability problem is PSPACE-complete; it has a very simple syntax, with only one metric operator; despite this, it is expressively equivalent with other very interesting logics, and in particular with the Quantitative Monadic Logic of Order (QMLO), and with the Metric Interval Temporal Logic (MITL). In fact, a translation has been defined that, from a QTL formula, produces an equivalent QMLO (resp. MITL) formula, and vice-versa. Since QMLO can be used to provide semantics to a variety of existing metric temporal logics, our approach can be used in principle to decide the satisfiability of a wide range of logics, including for example the popular MITL.

More precisely, in this paper we introduce a linear satisfiability-preserving translation from QTL formulae to formulae of CLTL-over-clocks (CLTL-oc), a decidable logic [BRS] whose satisfiability problem is also PSPACE-complete, for which it is possible to define a decision procedure based on Satisfiability Modulo Theories (SMT) solving techniques that are implemented in a variety of tools (such as [Mic]). This is the basis for a prototype tool, available from [qtl].

Although QTL is decidable over unrestricted models, we will focus on models that are *finitely variable*, 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 [FMMR12]) which do not have much practical interest.

The paper is organized as follows: Sect. 2 defines QTL and CLTL-oc, and Sect. 3 defines a reduction from the former to the latter; Sect. 4 shows that the translation is satisfiability-preserving, and discusses its complexity. Sect. 5 presents some experimental results carried out with our prototype tool. Sect. 6 concludes, describing also tool support. *All proofs can be found in the extended version of this paper that is available from the tool website [qtl].*

2 Languages

Let AP be a finite set of atomic propositions. The syntax of (well-formed) QTL formulae over AP is defined by the grammar (where $p \in AP$):

$$\phi := p \mid \phi \wedge \phi \mid \neg\phi \mid \phi \mathbf{U}_{(0,\infty)} \phi \mid \mathbf{F}_{(0,1)} \phi \mid \phi \mathbf{S}_{(0,\infty)} \phi \mid \mathbf{P}_{(0,1)} \phi.$$

The semantics of QTL may be defined with respect to a generic linear order, but in what follows we will focus on the nonnegative real line, i.e., the linear order $(\mathbb{R}_{\geq 0}, <)$. A *structure* M for QTL over alphabet AP is a pair $M = \langle \mathbb{R}_{\geq 0}, \mathcal{B}^M \rangle$, where \mathcal{B}^M is a valuation mapping every propositional variable $p \in AP$ to a set $\mathcal{B}^M(p) \subseteq \mathbb{R}_{\geq 0}$. Hence, a structure may be considered as providing continuous-time Boolean *signals* over the set AP . *Satisfaction* of a QTL formula over

$$\begin{aligned}
M, t &\models p \Leftrightarrow t \in \mathcal{B}^M(p) \\
M, t &\models \neg\phi \Leftrightarrow M, t \not\models \phi \\
M, t &\models \phi \wedge \psi \Leftrightarrow M, t \models \phi \text{ and } M, t \models \psi \\
M, t &\models \phi \mathbf{U}_{(0,\infty)} \psi \Leftrightarrow \exists t' > t, M, t' \models \psi \text{ and } \forall t'', t < t'' < t', M, t'' \models \phi \\
M, t &\models \mathbf{F}_{(0,1)} \phi \Leftrightarrow \exists t', t < t' < t + 1, M, t' \models \phi \\
M, t &\models \phi \mathbf{S}_{(0,\infty)} \psi \Leftrightarrow \exists t' < t, M, t' \models \psi \text{ and } \forall t'', t' < t'' < t, M, t'' \models \phi \\
M, t &\models \mathbf{P}_{(0,1)} \phi \Leftrightarrow \exists t', t - 1 < t' < t, M, t' \models \phi.
\end{aligned}$$

Table 1: Semantics of QTL.

M at a point $t \in \mathbb{R}_{\geq 0}$ is a relation \models defined inductively as in Table 1. Given a QTL formula ϕ , we indicate by $sub(\phi)$ the set of all subformulae occurring in ϕ .

In this paper, we will assume signals to have *finite variability*, i.e., in any bounded time interval there can only be a finite number of changes. Nevertheless, the following result holds.

Theorem 1 ([HR05]) *Satisfiability of QTL over $(\mathbb{R}_{\geq 0}, <)$ is PSPACE-complete, even without the finite variability assumption.*

Constraint LTL (CLTL [DD07, BFRS11]) formulae are defined with respect to a finite set V of variables and a structure $\mathcal{D} = (D, \mathcal{R})$ where D is a specific domain of interpretation for variables and constants and \mathcal{R} is a family of relations on D , with the set AP of atomic propositions being the set \mathcal{R}_0 of 0-ary relations. An *atomic constraint* is a term of the form $R(x_1, \dots, x_n)$, where R is an n -ary relation of \mathcal{R} on D and $x_1, \dots, x_n \in V$. A *valuation* is a mapping $v : V \rightarrow D$. A constraint is *satisfied* by v , written $v \models_{\mathcal{D}} R(x_1, \dots, x_n)$, if $(v(x_1), \dots, v(x_n)) \in R$.

Temporal terms α are defined by the syntax $\alpha := c \mid x \mid \mathbf{X}\alpha$, where c is a constant in D and $x \in V$. CLTL formulae are defined as follows:

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

where α_i 's are temporal terms, $R \in \mathcal{R}$, \mathbf{X} , \mathbf{Y} , \mathbf{U} and \mathbf{S} are the usual “next”, “previous”, “until” and “since” operators of LTL, with the same meaning. Operator \mathbf{X} is similar to \mathbf{X} , but it only applies to temporal terms, with the meaning that $\mathbf{X}\alpha$ is the *value* of temporal term α in the next time instant. Operators “globally” \mathbf{G} and “release” \mathbf{R} are introduced as customary as abbreviations: $\phi_1 \mathbf{R} \phi_2 = \neg(\neg\phi_1 \mathbf{U} \neg\phi_2)$, $\mathbf{G}(\phi) = \perp \mathbf{R} \phi$.

The *depth* $|\alpha|$ of a temporal term is the total amount of temporal shift needed in evaluating α : $|\alpha| = 0$ when α is a variable, and $|\mathbf{X}\alpha| = |\alpha| + 1$. 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 $\pi : \mathbb{N} \rightarrow \wp(AP)$ and $\sigma : \mathbb{N} \times V \rightarrow D$, which define a subset of AP and the value of variables for each element of \mathbb{N} . The value of terms is defined with respect to σ as follows:

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

$$\begin{aligned}
 (\pi, \sigma), i \models p &\Leftrightarrow p \in \pi(i) \text{ for } p \in AP \\
 (\pi, \sigma), i \models R(\alpha_1, \dots, \alpha_n) &\Leftrightarrow (\sigma(i + |\alpha_1|, x_{\alpha_1}), \dots, \sigma(i + |\alpha_n|, x_{\alpha_n})) \in R \\
 (\pi, \sigma), i \models \mathbf{X}(\phi) &\Leftrightarrow (\pi, \sigma), i + 1 \models \phi \\
 (\pi, \sigma), i \models \mathbf{Y}(\phi) &\Leftrightarrow (\pi, \sigma), i - 1 \models \phi \wedge i > 0 \\
 (\pi, \sigma), i \models \phi \mathbf{U} \psi &\Leftrightarrow \exists j \geq i : (\pi, \sigma), j \models \psi \wedge \forall i \leq n < j, (\pi, \sigma), n \models \phi \\
 (\pi, \sigma), i \models \phi \mathbf{S} \psi &\Leftrightarrow \exists 0 \leq j \leq i : (\pi, \sigma), j \models \psi \wedge \forall i \leq n < j, (\pi, \sigma), n \models \phi
 \end{aligned}$$

Table 2: Semantics of CLTL (propositional connectives are omitted for brevity).

assuming that x_α is the variable in V occurring in term α . The semantics of a CLTL formula ϕ at instant $i \geq 0$ over a pair (π, σ) is defined as in Table 2, where x_{α_i} is the variable that appears in temporal term α_i , and $R \in \mathcal{R} \setminus \mathcal{R}_0$ (recall that $\mathcal{R}_0 = AP$). A formula CLTL ϕ is *satisfiable* if there exists a pair (π, σ) such that $(\pi, \sigma), 0 \models \phi$; in this case, we say that (π, σ) is a *model* of ϕ .

In this paper, we restrict the set of models where variables in V are evaluated as *clocks*. A clock “measures” the time elapsed since its last “reset” (i.e., the variable was equal to 0). Each position $i \in \mathbb{N}$ is associated with a “time delay” $\delta(i)$, where $\delta(i) > 0$ for all i , corresponding to the “time elapsed” between the current position i and the next one $i + 1$. For a clock x_α ,

$$\sigma(i + 1, x_\alpha) = \begin{cases} \sigma(i, x_\alpha) + \delta(i), & \text{time elapsing} \\ 0 & \text{reset } x_\alpha. \end{cases}$$

The set \mathcal{R} is restricted to $\{<, =\}$ because CLTL-oc formulae need only to measure the time elapsing among events, as later explained. Under these two restrictions, CLTL-oc is decidable [BRS], and an effective decision procedure can be devised by encoding CLTL-oc formulae into formulae in the decidable theory of Quantifier-free Uninterpreted Functions with Equality combined with Linear Real Arithmetic (QF-EUF \cup LRA), which is solved by SMT solvers such as, for example, Z3 [Mic]. A prototype solver for CLTL-oc formulae is available as part of the Zot tool [ae2].

QTL is closely related to other metric temporal logics, and in particular QMLO [HR05] and the popular MITL [AFH96], through the following result.

Theorem 2 ([HR05]) *QMLO, QTL and MITL are expressively equivalent.*

Hence, a satisfiability-preserving translation of QTL formulae into CLTL-oc ones can be the basis for an effective decision procedure to solve the satisfiability (over finitely-variable behaviors) of all above-mentioned logics.

3 Reduction of QTL to CLTL-over-clocks

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




Figure 1: Example of QTL signal and a corresponding CLTL-oc model (clocks not shown).

is discrete. CLTL-oc variables behaving as clocks represent time progress, while discrete positions in CLTL-oc models represent, for each subformula occurring in QTL formula ϕ , whether a change of truth value (an “event”) occurs or not for the subformula at that point. Time progress between two discrete points is measured by CLTL-oc clocks; between events, the truth value of formulae is stable (i.e., there is no change). In every (discrete) position CLTL-oc models embed, through suitable fresh propositional letters (\uparrow and \leftarrow), the information defining the truth value of all the subformulae occurring in QTL formula ϕ and, through clock variables, the information about the time progress between two consecutive changing points. Then, every position in a CLTL-oc model captures the configuration of one of the intervals in which the continuous-time signal is partitioned by considering the QTL “events”. Therefore, our reduction defines, by means of CLTL-oc formulae, the semantics of every subformula occurring in ϕ . Fig. 1 shows an example of QTL signal and a corresponding CLTL-oc model.

Consider a QTL formula ϕ . For each subformula θ of ϕ we introduce two predicates, \uparrow_{θ} and \leftarrow_{θ} , which represent the value of θ in, respectively, the first instant and the rest of the interval between two events (hence, \uparrow_{θ} represents the value of θ exactly when the event occurs). We also introduce two clocks, z_{θ}^0 and z_{θ}^1 , which measure the time elapsed since the last two “events”.

Let $\theta \in \text{sub}(\phi)$. We say that the event “ θ becomes true” e_{θ}^u occurs at instant $t \geq 0$ of signal M when θ holds right after t , but not before it, or t is the origin:

$$\exists \varepsilon > 0, \forall t' \in (t, t + \varepsilon) \text{ it is } M, t' \models \theta \text{ and either } t = 0 \text{ or } \exists \varepsilon' > 0, \forall t' \in (t - \varepsilon', t) \text{ it is } M, t' \models \neg \theta.$$

The opposite event “ θ becomes false” e_{θ}^d is simply given by the property above with $\neg \theta$ instead of θ . QTL events e_{θ}^u and e_{θ}^d are represented in the CLTL-oc formula through combinations of the basic predicates \uparrow_{θ} and \leftarrow_{θ} that are abbreviated by \sqcup_{θ} and \sqcap_{θ} , respectively, whose definitions are shown in Table 3.

We do not impose any restrictions on signals other than they be finitely variable. In particular, subformulae θ can have *singularities*, i.e., instants in which the value of θ is different than in their neighborhood. More precisely, we say that a formula θ has an “up-singularity” s_{θ}^u in instant t if the following holds:

$$t > 0, M, t \models \theta \text{ and } \exists \varepsilon > 0 \text{ s.t. } \forall t' \neq t \in (t - \varepsilon, t + \varepsilon) \text{ it is } M, t' \models \neg \theta.$$

\uparrow_ξ	= ξ holds in the first instant of the current interval		
$\overleftarrow{\xi}$	= ξ holds in the current interval (except possibly for its first instant)		
\Downarrow_ξ	= $\neg \mathbf{Y}(\overleftarrow{\xi}) \wedge \overleftarrow{\xi}$	\Downarrow_ξ	= $\mathbf{Y}(\neg \overleftarrow{\xi}) \wedge \uparrow_\xi \wedge \neg \overleftarrow{\xi}$
\Uparrow_ξ	= $\neg \mathbf{Y}(\neg \overleftarrow{\xi}) \wedge \neg \overleftarrow{\xi}$	\Uparrow_ξ	= $\mathbf{Y}(\overleftarrow{\xi}) \wedge \neg \uparrow_\xi \wedge \overleftarrow{\xi}$
$\overleftarrow{\downarrow}_\xi$	= $\Uparrow_\xi \vee \Downarrow_\xi \vee (\text{orig} \wedge \neg \uparrow_\xi)$	$\overleftarrow{\downarrow}_\xi$	= $\Uparrow_\xi \vee \Downarrow_\xi$
$\overleftarrow{\uparrow}_\xi$	= $\Downarrow_\xi \vee \Uparrow_\xi \vee (\text{orig} \wedge \uparrow_\xi)$	$\overleftarrow{\uparrow}_\xi$	= $\Downarrow_\xi \vee \Uparrow_\xi$
$\overleftarrow{\xi}$	= $\uparrow_\xi \wedge \overleftarrow{\xi}$	$\overleftarrow{\xi}$	= $\neg \uparrow_\xi \wedge \neg \overleftarrow{\xi}$
orig	= $\neg \mathbf{Y}(\top)$		

Table 3: CLTL-oc predicates and abbreviations used in the encoding. Note that $\mathbf{Y}(\overleftarrow{\xi})$ and $\mathbf{Y}(\neg \overleftarrow{\xi})$ are false in the origin, no matter ξ , and elsewhere $\neg \mathbf{Y}(\neg \overleftarrow{\xi}) \equiv \mathbf{Y}(\overleftarrow{\xi})$; hence, \Downarrow_ξ holds in 0 if, and only if, $\overleftarrow{\xi}$ holds there, \Downarrow_ξ does not hold in 0, and so on.

We say that θ has a “down-singularity” s_θ^d if the formula above holds with $\neg\theta$ instead of θ . Note that singularities do not occur in the origin. In CLTL-oc, we represent up- and down-singularities with combinations of basic propositions abbreviated by \Downarrow_θ and \Uparrow_θ , respectively, as shown in Table 3.

Table 3 summarizes the CLTL-oc predicates used here. In a nutshell, $\overleftarrow{\downarrow}_\xi$ (resp. $\overleftarrow{\uparrow}_\xi$) indicates that formula ξ held (resp. did not hold) in an interval before the current one, and now it switches; the switch can be singular (in which case ξ immediately takes the same value it held before now), or not, in which case ξ stays false (resp. true) for some time after the switch. Formula $\overleftarrow{\downarrow}_\xi$ (resp. $\overleftarrow{\uparrow}_\xi$), instead, holds if ξ becomes true (resp. false) in the current instant, and it holds in an interval after now. Also, formula $\overleftarrow{\xi}$ (resp. $\overleftarrow{\xi}$) states that ξ is true (resp. false) throughout the current interval. In addition, we abbreviate by orig formula $\neg \mathbf{Y}(\top)$, which holds only in 0.

In the rest of this section we define the translation from QTL to CLTL-oc formulae which is the main contribution of this paper. First, Section 3.1 introduces a set of general formulae, which are written for any subformula θ of ϕ , defining constraints that guarantee that clock resets occur at suitable points. Then, in Section 3.2, we provide the operator-specific CLTL-oc formulae that capture the semantics of QTL connectives and temporal operators.

3.1 General Constraints on Clocks and Events

This section describes the behavior of clocks and events. We introduce clocks z_θ^0 and z_θ^1 for each subformula θ of ϕ to measure the time elapsing between two consecutive events of θ . In each discrete position of a CLTL-oc model, the value of z_θ^0 and z_θ^1 is, intuitively, the time elapsed since the last two events of θ , which is set to 0 (reset) only when an event (of θ) occurs. Resets of z_θ^0

and z_θ^1 alternate because, when one of the two clocks is reset to start measuring the time elapsing from the current event, the time elapsed since the previous event (which is needed in CLTL-oc formulae to model the semantics of QTL modalities) is measured by the other clock. In other words, one can not “read” the value of a clock and, at the same time, reset it to start measuring the elapsed time anew.

For any $\theta \in \text{sub}(\phi)$, the following CLTL-oc formula holds at position 0, simply stating that in 0 the z_θ^0 clock of every subformula θ is reset (while z_θ^1 can have any value):

$$z_\theta^0 = 0 \quad (1)$$

The other formulae of this section must hold at each discrete instant; for simplicity, the globally operator \mathbf{G} is inserted explicitly only at the end of the section.

Whenever subformula θ switches its value (it becomes true or false, possibly in a singular way), one of its associated clocks z_θ^0 and z_θ^1 is reset:

$$\overset{\theta}{\uparrow} \vee \overset{\theta}{\downarrow} \Leftrightarrow z_\theta^0 = 0 \vee z_\theta^1 = 0. \quad (2)$$

The clocks associated with a subformula θ are reset in an alternate way: between any two resets of clock z_θ^0 there must be a reset of clock z_θ^1 , and vice-versa:

$$\bigwedge_{i \in \{0,1\}} (z_\theta^i = 0) \Rightarrow \mathbf{X} \left((z_\theta^{(i+1) \bmod 2} = 0) \mathbf{R}(z_\theta^i \neq 0) \right). \quad (3)$$

In the following, genconstr_θ denotes the formula $(1) \wedge \mathbf{G}((2) \wedge (3))$.

3.2 Semantics of QTL temporal modalities

This section presents the definition of $m(\theta)$, the translation of every subformula θ of a QTL formula into a suitable CLTL-oc formula encoding its semantics. Essentially, $m(\theta)$ describes how θ becomes true and false depending on the value of its own subformulae.

- $\theta = \neg\psi$: The predicates related to θ are exactly the opposite ones of ψ , so $m(\theta)$ is the following:

$$m(\theta) = (\uparrow_\theta \Leftrightarrow \neg \uparrow_\psi) \wedge (\overset{\leftarrow}{\theta} \Leftrightarrow \neg \overset{\leftarrow}{\psi}). \quad (4)$$

- $\theta = \gamma \wedge \psi$: The semantics of $\gamma \wedge \psi$ is simply the conjunction of the basic predicates for γ and ψ .

$$m(\theta) = (\uparrow_\theta \Leftrightarrow \uparrow_\gamma \wedge \uparrow_\psi) \wedge (\overset{\leftarrow}{\theta} \Leftrightarrow \overset{\leftarrow}{\gamma} \wedge \overset{\leftarrow}{\psi}) \quad (5)$$

- $\theta = \gamma \mathbf{U}_{(0,+\infty)} \psi$: The following lemma holds for formulae of this form.

Lemma 1 *Let $\theta = \gamma \mathbf{U}_{(0,+\infty)} \psi$ and M be a non-Zeno signal. For each $t \in \mathbb{R}_+$ there is $\varepsilon \in \mathbb{R}_{>0}$ such that $M, t \models \theta$ if, and only if, for all $t' \in (t, t + \varepsilon]$ it is $M, t' \models \theta$.*

Then, \mathbf{U} formulae can not have singularity points, as they would violate Lemma 1. In addition, when a \mathbf{U} formula changes its value, it must do so in a left-closed manner (i.e., the value at the change point is the same as the one after the change point) or, again, Lemma 1 is violated. Then, we have (6) below.

$$m(\theta) = \left(\uparrow_{\theta} \Leftrightarrow \overleftarrow{\theta} \right) \wedge \left(\overleftarrow{\theta} \Leftrightarrow \overleftarrow{\gamma} \wedge \left(\overleftarrow{\psi} \vee \mathbf{X} \left(\overleftarrow{\gamma} \mathbf{U} \left((\overleftarrow{\gamma} \wedge \overleftarrow{\psi}) \vee \uparrow_{\psi} \right) \right) \right) \right) \quad (6)$$

In particular, the second conjunct of Formula (6) states that θ holds in an interval if, and only if, either both ψ and γ hold in it, or there is a future interval in which ψ holds (either throughout the interval, or in its first instant), and γ holds throughout all intervals (including their first instants) in between.

- $\theta = \mathbf{F}_{(0,1)}\gamma$: For formulae $\mathbf{F}_{(0,1)}\gamma$ we have the following result.

Lemma 2 *Let $\theta = \mathbf{F}_{(0,1)}\gamma$ be a QTL formula. If $M, t \models \theta$ then there is $\varepsilon \in \mathbb{R}_{>0}$ 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 2, an up-singularity \lceil_{θ} can never occur for a formula of the form $\mathbf{F}_{(0,1)}\gamma$. In addition, if θ holds at the beginning of an interval (i.e., \uparrow_{θ} 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:

$$\uparrow_{\theta} \Rightarrow \overleftarrow{\theta} \wedge (\mathbf{Y}(\overleftarrow{\theta}) \vee \text{orig}) \quad (7)$$

Formula (8) states that, when θ becomes true with a raising edge \lceil_{θ} , in an instant other than the origin, a clock z_{θ}^j is reset, and \uparrow_{γ} will eventually be true after 1 instant; if θ becomes true in the origin, then either it does so in a left-closed manner, and γ becomes true before clock z_{θ}^0 becomes 1, or it becomes true in a left-open manner, and γ becomes true exactly at 1. Fig. 2(a) gives a graphical depiction of one of the conditions for having a raising edge in $t > 0$.

$$\lceil_{\theta} \Leftrightarrow \left(\begin{array}{l} \text{orig} \wedge \left(\uparrow_{\theta} \wedge \left(\lceil_{\gamma} \vee \mathbf{X} \left(z_{\theta}^0 > 0 \mathbf{U} \left(\uparrow_{\gamma} \wedge (0 < z_{\theta}^0 < 1) \right) \right) \right) \right) \vee \\ \neg \uparrow_{\theta} \wedge \lceil_{\gamma} \wedge \mathbf{X} \left((z_{\theta}^0 > 0 \wedge \lceil_{\gamma}) \mathbf{U} \left(\uparrow_{\gamma} \wedge z_{\theta}^0 = 1 \right) \right) \end{array} \right) \vee \\ \left(\neg \text{orig} \wedge \neg \uparrow_{\theta} \wedge \bigvee_{j \in \{0,1\}} \left(z_{\theta}^j = 0 \wedge \mathbf{X} \left(z_{\theta}^j > 0 \mathbf{U} \left(\uparrow_{\gamma} \wedge z_{\theta}^j = 1 \wedge \bigvee_{i \in \{0,1\}} z_{\gamma}^i > 1 \right) \right) \right) \right) \quad (8)$$

We also add a constraint, which is captured by Formula (9), which states that, if γ becomes true in an instant t , and it was false in the interval of length 1 preceding t , then in t one of the clocks associated with θ has value 1, since $\mathbf{F}_{(0,1)}\gamma$ started holding 1 time unit before t .

$$\left(\uparrow_{\gamma} \wedge \bigvee_{i \in \{0,1\}} z_{\gamma}^i \geq 1 \right) \Rightarrow \bigvee_{j \in \{0,1\}} z_{\theta}^j = 1 \quad (9)$$



Figure 2: Depiction of some conditions for raising and falling edges in metric operators.

When θ becomes false with either a falling edge (\downarrow_{θ}) or in a singular manner (\Uparrow_{θ}), γ becomes false, so a clock z_{γ}^i is reset. If θ becomes false with a falling edge (10), then γ can not become true again as long as the clock that is reset with \downarrow_{γ} is ≤ 1 . If θ becomes false in a singular manner (11), instead, γ must become true again exactly when the clock that is reset with \downarrow_{γ} is 1.

$$\downarrow_{\theta} \Leftrightarrow \downarrow_{\gamma} \wedge \neg \mathbf{X} \left(\downarrow_{\gamma} \mathbf{U}(\uparrow_{\gamma} \wedge \bigvee_{i \in \{0,1\}} 0 < z_{\gamma}^i \leq 1) \right) \quad (10)$$

$$\Uparrow_{\theta} \Leftrightarrow \downarrow_{\gamma} \wedge \mathbf{X} \left(\downarrow_{\gamma} \mathbf{U}(\uparrow_{\gamma} \wedge \bigvee_{i \in \{0,1\}} z_{\gamma}^i = 1) \right) \wedge \neg \text{orig} \quad (11)$$

Then, for $\theta = \mathbf{F}_{(0,1)}\gamma$, $m(\theta)$ is (7) \wedge (8) \wedge (10) \wedge (11).

Case $\theta = \gamma \mathbf{S}_{(0,+\infty)} \psi$

In this case, we have a result that is similar to Lemma 1:

Lemma 3 *If $\theta = \gamma \mathbf{S}_{(0,+\infty)} \psi$ and M is a non-Zeno signal, then, for each $t \in \mathbb{R}_+$ there is $\varepsilon \in \mathbb{R}_{>0}$ such that $M, t \models \theta$ if, and only if, for all $t' \in [t - \varepsilon, t)$ it is also $M, t' \models \theta$.*

Note that in $t = 0$ $\gamma \mathbf{S}_{(0,+\infty)} \psi$ is false, and, for any $\varepsilon \in \mathbb{R}_{>0}$, $[-\varepsilon, 0)$ is not an interval of \mathbb{R}_+ , so the proposition is trivially true.

Then, \mathbf{S} formulae can not have singularity points, as they would violate Lemma 3. 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). In the origin, instead, θ is false. Then, we have

$$m(\theta) = \left(\uparrow_{\theta} \Leftrightarrow \mathbf{Y}(\overleftarrow{\theta}) \right) \wedge \left(\overleftarrow{\theta} \Leftrightarrow \overleftarrow{\gamma} \mathbf{S}(\left(\uparrow_{\psi} \vee \overleftarrow{\psi} \right) \wedge \overleftarrow{\gamma}) \right) \quad (12)$$

Case $\theta = \mathbf{P}_{(0,1)}\gamma$

For formulae $\mathbf{P}_{(0,1)}\gamma$ we have the following result.

Lemma 4 *Let $\theta = \mathbf{P}_{(0,1)}\gamma$ be a QTL formula; if θ holds for a signal M in an instant t (i.e., $M, t \models \theta$), then there is $\varepsilon \in \mathbb{R}_{>0}$ such that, for all $t' \in [t - \varepsilon, t + \varepsilon]$ it is also $M, t' \models \theta$.*

Note that $\mathbf{P}_{(0,1)}\gamma$ is false in $t = 0$, no matter γ . As for **F** formulae, Lemma 4 implies that $\lrcorner\lrcorner\theta$ can never occur for θ . In addition, by Lemma 4, if θ holds in the first instant of an interval t (i.e., $\uparrow\theta$), it must also hold in the intervals before and after t . Then, the following constraint holds:

$$\uparrow\theta \Rightarrow \overset{\leftarrow}{\theta} \wedge \mathbf{Y}\left(\overset{\leftarrow}{\theta}\right) \quad (13)$$

Formula (14) states that for θ to become true with a raising edge in t , γ must also become true (possibly in a singular manner). This is sufficient if $t = 0$. If $t > 0$, there are two cases: either γ was never true before t (so it was false in the origin and it stayed so), or the last change point of γ before t was before $t - 1$, so the clock associated with γ that is not reset in t is > 1 .

$$\lrcorner\theta \Leftrightarrow \overset{\gamma}{\uparrow} \wedge \left(\text{orig} \vee \mathbf{Y}\left(\overset{\gamma}{\downarrow} \mathbf{S}(\text{orig} \wedge \underline{\gamma})\right) \vee \bigvee_{i \in \{0,1\}} z_{\gamma}^i > 1 \right) \quad (14)$$

Formula (15) states that θ has a falling edge in t if and only if either $t = 0$ and there is ε such that γ is false in $[0, \varepsilon)$, or the last time γ became true was at $t - 1$. This corresponds to the condition (depicted in Fig. 2(b)) that there is a z_{γ}^i that is 1 in t , and the last time γ had a change point it was $z_{\gamma}^i = 0$ and γ became false. γ can not become true in t , or θ would not have a falling edge; if γ becomes true in t , then θ has a down-singularity, as specified by Formula (16).

$$\lrcorner\theta \Leftrightarrow \bigvee_{i \in \{0,1\}} \left(z_{\gamma}^i = 1 \wedge \left(\overset{\gamma}{\downarrow} \mathbf{S} \left(\overset{\gamma}{\downarrow} \wedge z_{\gamma}^i = 0 \right) \right) \right) \vee (\text{orig} \wedge \underline{\gamma}) \quad (15)$$

$$\lrcorner\theta \Leftrightarrow \overset{\gamma}{\uparrow} \wedge \bigvee_{i \in \{0,1\}} \left(z_{\gamma}^i = 1 \wedge \mathbf{Y} \left(\overset{\gamma}{\downarrow} \mathbf{S} \left(\overset{\gamma}{\downarrow} \wedge z_{\gamma}^i = 0 \wedge \neg(\text{orig} \wedge \underline{\gamma}) \right) \right) \right) \quad (16)$$

Finally, we introduce the analogous for the eventuality in the past of Formula (9). More precisely, Formula (17) specifies that if γ becomes false and there are no events associated with γ for at least 1 time unit, the CLTL-oc model includes a position in which the clock that is reset with the falling edge of γ hits value 1. Formula (17) is necessary to make sure that, if γ becomes false (and it does not become true again for 1 time unit, hence θ must also become false after 1), eventually the right hand side of Formulae (15) and (16) holds.

$$\bigwedge_{i \in \{0,1\}} \overset{\gamma}{\downarrow} \wedge z_{\gamma}^i = 0 \Rightarrow (z_{\gamma}^i < 1) \mathbf{U} \left(z_{\gamma}^i = 1 \vee (\overset{\gamma}{\uparrow} \wedge 0 < z_{\gamma}^i < 1) \right) \quad (17)$$

Then, for $\theta = \mathbf{P}_{(0,1)}(\gamma)$, $m(\theta)$ is (13) \wedge (14) \wedge (15) \wedge (16) \wedge (17).

Finally, QTL formula ϕ is initially satisfiable if, and only if, it holds in the first instant of the interval starting at 0, i.e., $\text{init}_\phi = \uparrow_\phi$. Then, for a QTL formula ϕ , the corresponding CLTL-oc formula ϕ_{CLTL} is:

$$\phi_{\text{CLTL}} = \text{init}_\phi \wedge \bigwedge_{\theta \in \text{sub}(\phi)} (\text{genconstr}_\theta \wedge \mathbf{G}(m(\theta))). \quad (18)$$

The next section shows the correctness of the translation.

4 Correctness and complexity of the reduction

To complete the results of this paper, we need to show that a QTL formula ϕ is satisfiable if, and only if, there exists a pair (π, σ) that satisfies ϕ_{CLTL} defined by (18).

First of all, we define a correspondence between QTL signals and CLTL-oc interpretations. Let us consider a finitely variable signal M that is an interpretation for a QTL formula θ ; we call $r_\theta(M)$ the set of CLTL-oc interpretations (π, σ) built according to the rules presented below.

Since M is finitely variable, the set of “events” in M for formula θ is denumerable. Let $T = \{t_k\}_{k \in \mathbb{N}} \subset \mathbb{R}_+$ be a denumerable set of time instants such that $t_k < t_j \Leftrightarrow k < j$, for all $t' \in \mathbb{R}_+$ there is $t_k \in T$ such that $t_k > t'$, and if t is an instant when at least one event for θ occurs in M , then $t \in T$. In the following we say that a clock v is reset at position k when $\sigma(k, v) = 0$.

If one event among $e_\theta^u, e_\theta^d, s_\theta^u$ or s_θ^d occurs at $t_k \in T$, the event marker captured by the corresponding formula $\lceil_\theta, \lfloor_\theta, \lceil\lceil_\theta, \lceil\lceil\lceil_\theta$ holds in $\pi(k)$; that is, if $M, t_k \models e_\theta^u$, then \lceil_θ holds in $\pi(k)$ (hence $\lceil\lceil_\theta \notin \pi(k-1)$, $\lceil\lceil_\theta \in \pi(k)$), and so on. In addition, if $M, t_k \models e_\theta^u$ and $M, t_k \models \theta$ (resp. $M, t_k \not\models \theta$), then $\uparrow_\theta \in \pi(k)$ (resp. $\uparrow_\theta \notin \pi(k)$); similarly for the falling edge. By the definition of events given in Sect. 3, θ has an event in $t = 0$, so $t_0 = 0$. If in $t_k \in T$ no events for θ occur, then none of $\{\lceil_\theta, \lfloor_\theta, \lceil\lceil_\theta, \lceil\lceil\lceil_\theta\}$ holds in $\pi(k)$ (so $\lceil\lceil_\theta \in \pi(k-1)$ iff $\uparrow_\theta \in \pi(k)$).

For each $t_k \in T$ where an event for θ occurs, either z_θ^0 or z_θ^1 is reset at k . z_θ^0 is reset in 0; after 0, clocks are reset modulo 2, i.e., if $\sigma(k, z_\theta^i) = 0$, and $\sigma(k', z_\theta^i) = 0$, where $i \in \{0, 1\}$ and $k' > k$, then there is a $k < j < k'$ s.t. $\sigma(j, z_\theta^{(i+1) \bmod 2}) = 0$. For each clock z_θ^i it is $\sigma(k+1, z_\theta^i) = \sigma(k, z_\theta^i) + t_{k+1} - t_k$ unless z_θ^i is reset.

Note that for a given signal M there is more than one possible compatible set $T = \{t_k\}_{k \in \mathbb{N}}$, and each one corresponds to a different CLTL-oc interpretation (for example, a signal in which $AP = \{p\}$ and p is always true is compatible with a set in which $t_k = k$, with one in which $t_k = 2k$, and so on). However, one can show that if two signals $M_1 \neq M_2$ differ for θ in at least one instant $t \in \mathbb{R}_+$, $r_\theta(M_1) \cap r_\theta(M_2) = \emptyset$. Then, given a CLTL-oc interpretation (π, σ) , there is at most one signal M such that $(\pi, \sigma) \in r_\theta(M)$; hence, we define $r_\theta^{-1}((\pi, \sigma))$ as the function that, given a CLTL-oc interpretation, returns the corresponding QTL signal, if any.

Consider a set \mathcal{F} of formulae; with an abuse of notation denote with $r_{\mathcal{F}}(M)$ the set of CLTL-oc interpretations built as above, but considering every event related to the formulae in \mathcal{F} . Given a formula ϕ , we focus on $r_{\text{sub}(\phi)}(M)$.

Not all CLTL-oc interpretations (π, σ) represent QTL signals, so there are pairs (π, σ) such that $r_\theta^{-1}((\pi, \sigma)) = \perp$ (where \perp represents that the function is not defined). However, we have the following results.

Lemma 5 *Let θ be a QTL formula and M a signal. For all interpretations (π, σ) such that $(\pi, \sigma) \in r_\theta(M)$ it is $(\pi, \sigma), 0 \models \text{genconstr}_\theta$.*

Lemma 6 *Let θ be a QTL formula and (π, σ) a CLTL-oc interpretation over $\uparrow_\theta, \overleftarrow{\theta}$ where time diverges (i.e., where $\sum_{i \in \mathbb{N}} \delta(i) = \infty$). Then, there is exactly one signal M such that $(\pi, \sigma) \in r_\theta(M)$.*

From the above results we have that, given a QTL formula ϕ , formula $\bigwedge_{\theta \in \text{sub}(\phi)} \text{genconstr}_\theta$ captures exactly all CLTL-oc interpretations such that $r_{\text{sub}(\phi)}^{-1}((\pi, \sigma)) \neq \perp$. Then, we have the following result.

Lemma 7 *Let M be a signal, and ϕ a QTL formula. For any $(\pi, \sigma) \in r_{\text{sub}(\phi)}(M)$ it is $(\pi, \sigma), 0 \models \bigwedge_{\theta \in \text{sub}(\phi)} \text{genconstr}_\theta$ and for all $k \in \mathbb{N}, \theta \in \text{sub}(\phi)$ it is $(\pi, \sigma), k \models m(\theta)$. Conversely, if $(\pi, \sigma), 0 \models \bigwedge_{\theta \in \text{sub}(\phi)} \text{genconstr}_\theta \wedge \mathbf{G}(m(\theta))$ and $M = r_{\text{sub}(\phi)}^{-1}((\pi, \sigma))$, then $(\pi, \sigma), k \models \neg_\phi$ if, and only if, $M, t_k \models e_\phi^u$ (similarly for the other events), and $\uparrow_\phi \in \pi(k)$ if, and only if, $M, t_k \models \phi$.*

Finally, from Lemma 7 the following theorem descends by observing that signal M is model for ϕ if, and only if, $M, 0 \models \phi$, which means that in $0 \uparrow_\phi$ holds.

Theorem 3 *Let ϕ be a QTL formula. ϕ is satisfiable if, and only if, ϕ_{CLTL} defined by (18) is satisfiable.*

Consider a QTL formula ϕ . The translation provided in Sect. 3 introduces, for each $\theta \in \text{sub}(\phi)$, 2 atomic propositions $\uparrow_\theta, \overleftarrow{\theta}$ and 2 variables z_θ^0, z_θ^1 . All CLTL-oc formulae $m(\theta)$ have fixed size. Hence, the size of Formula (18) linearly depends on the size of ϕ . [BRS] shows that satisfiability for a CLTL-oc formula ϕ_{CLTL} is PSPACE in the number of subformulae of ϕ_{CLTL} and the maximum constant occurring in it (which is 1 in the case of QTL). Then our translation preserves the PSPACE complexity of the satisfiability of QTL [HR05].

5 Some Experimental Results

The reduction of Sect. 3 is implemented in the `qtlsolver` tool, available from [qtl] and described in some further detail in [BRS]. The tool translates QTL into CLTL-oc, which can be checked for satisfiability by `ae2zot`, a plugin of the Zot bounded satisfiability checking tool available from [ae2].

The current implementation of `qtlsolver` supports various reductions. In particular, it implements a translation from a generalized version of QTL to CLTL-oc. This translation does not assume any special shape for signals, except that they be finitely variable; it natively supports operators $\mathbf{F}_{(0,b)}$ and $\mathbf{G}_{(0,b)}$ (and their past counterparts). These operators allow us to define concisely MITL operators [AFH96] $\mathbf{F}_{\langle a,b \rangle}$ and $\mathbf{G}_{\langle a,b \rangle}$ as abbreviations, where bounds can be either included or excluded. For instance, $\mathbf{G}_{(3,6)}(\phi)$ is equivalent to $\mathbf{G}_{(0,3)}(\mathbf{F}_{(0,3)}(\mathbf{G}_{(0,3)}(\phi)))$.

We used the `qtlsolver` tool to perform satisfiability checks on some examples (see also the

tool website [qtl]). Let us briefly introduce a pair of them, the first one taken from an LTL specification of [PMS12].

Consider a lamp controlled by two buttons, labeled ON and OFF respectively, which can not be pressed simultaneously. The lamp itself can be either on or off. When ON is pressed the lamp is immediately turned on, regardless of its current state, while if OFF is pushed then the lamp is immediately turned off, also regardless of its current state. However, to save energy there is also a timeout: after ON is pressed, the lamp will not stay on forever, but, if no more buttons are pressed, it will automatically turn off with a delay Δ , a positive real constant. Notice that, from this definition, it follows that by pressing the ON button before the timeout expiration then the timeout is extended by a new delay Δ .

We built a QTL specification of the timed lamp that uses atomic propositions on, off and l representing, respectively, events “push button ON” and “push button OFF” and the state “light is on”. We introduced constraints that specify that predicates on and off are constrained to be true only in isolated instants.

On this specification we have carried out three experiments: a check of the satisfiability of the specification, to show that it is consistent (*sat*); the (dis)proof of property “the light never stays on for more than Δ time units” (p_1); the proof of property “if at some point the light stays on for more than Δ time units, then there is an instant in which the on button is pressed, and then it is pressed again before Δ time units” (p_2).¹

The behavior of the timed lamp can be captured by the following QTL formula (we write \mathbf{G} for $\mathbf{G}_{[0,\infty)}$, and \mathbf{S} for $\mathbf{S}_{[0,\infty)}$):

$$\mathbf{G} \left((l \Leftrightarrow (\neg off \mathbf{S} on) \wedge \mathbf{P}_{[0,\Delta)}(on)) \wedge (on \Rightarrow \neg off) \right). \quad (19)$$

As mentioned above, we force on to hold only in isolated instants by adding the following QTL constraint (similarly for off):

$$\mathbf{G}(\neg(on \mathbf{U}_{(0,+\infty)} \top) \wedge \neg(off \mathbf{S}_{(0,+\infty)} \top)). \quad (20)$$

Properties p_1 and p_2 are captured by the following QTL formulae (where \mathbf{F} stands for $\mathbf{F}_{[0,+\infty)}$):

$$\mathbf{G}(\mathbf{F}_{[0,\Delta)}(\neg l)) \quad (21)$$

$$\mathbf{F}(\mathbf{G}_{[0,\Delta)}(l) \Rightarrow \mathbf{F}(on \wedge \mathbf{F}_{(0,\Delta)}(on))). \quad (22)$$

Table 4 reports the time and space required for the checks outlined above.² All bounded satisfiability checks have been performed using a bound $k = 20$. The first line of each row shows the total processing time (i.e., parsing and solving) and the time taken by the SMT-solver (both times in seconds). The second line reports the heap size (in Mbytes) required by Z3. The results of the checks are the following: the specification is satisfiable, property p_1 does not hold (the tool returns a counterexample), while property p_2 holds (“unsat” is returned).

Finally, we present a behavior that highlights some interesting features of the tool. The behavior is captured by the following formulae, which state that p and q only occur in isolated instants,

¹ In all experiments it is $\Delta = 5$.

² All tests have been done using the Common Lisp compiler SBCL 1.1.2 on a 2.13GHz Core2 Duo MacBook Air with MacOS X 10.7 and 4GB of RAM. The solver was z3 4.0.

Problem	Satisfiable?	Time (Total/SMT only)	Memory
sat	Yes	4.24/3.04	27.12
p1	Yes	17.2/14.86	63.5
p2	No	257.1/240.88	58.66

Table 4: Experimental results with the timed lamp, reporting Time (sec) and heap size (MB).

with p occurring exactly every 80 time units, and q occurring within 80 time units in the past from each p (origin excluded).

$$\mathbf{G} \left(\begin{array}{l} \mathbf{G}_{(0,80)}(\neg p) \Rightarrow \mathbf{G}_{(80,160)}(\neg p) \quad \wedge \\ (p \Rightarrow \mathbf{F}_{(0,160)}p) \quad \wedge \quad (q \Rightarrow (\neg q) \mathbf{U} \top) \end{array} \right) \wedge p \wedge \mathbf{G}_{(0,80)}(\neg p) \wedge \mathbf{G}_{(0,\infty)}(p \Rightarrow \mathbf{P}_{(0,80)}q) \quad (23)$$

In this case, the bound $k = 10$ is enough to prove that the formula is satisfiable and a model is produced in about 40 secs. In around the same time the solver shows that property $\mathbf{G}(p \Rightarrow \mathbf{F}_{(0,80)}(q))$ holds for model (23) (up to the considered bound), whereas property $\mathbf{G}(q \Rightarrow \mathbf{F}_{(0,80)}(q))$ does not.

Note that, in Formula (23), the constants involved in the temporal modalities are significantly larger than the bound k required to obtain a model satisfying the formula. In fact, any value is possible in principle for the increments of the clocks between two consecutive discrete instants, controlled by the (nondeterministic) variable δ . This highlights that the length of the intervals described by a CLTL-oc model is independent of the bound k , as long as this is big enough to capture all changepoints that are necessary to build a periodic sequence of clock regions.

6 Conclusions

This paper presents a satisfiability-preserving translation from QTL formulae to formulae of the CLTL-oc logic, which can be solved through SMT solvers. As formulae of other logics such as QMLO and MITL can be in turn translated into equivalent QTL formulae, our encoding can be the basis for an effective decision procedure for several interesting logics.

The encoding presented in this paper has been implemented in a prototype tool [qtl]. Preliminary experiments are promising as we were able to solve some simple, yet conceptually significant, temporal behaviors in a reasonable amount of time. All these examples can be realized by discrete CLTL-oc models of short length, even when the time constants are quite big (provided the ratio among them is small). The outcome of the procedure is not only sat/unsat, but also (when applicable) a concrete model satisfying the formula.

Acknowledgements: Work supported by the Programme IDEAS-ERC, Project 227977-SMScom.

Bibliography

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

- [ae2] Zot: a Bounded Satisfiability Checker. available from zot.googlecode.com.
- [AFH96] R. Alur, T. Feder, T. A. Henzinger. The Benefits of Relaxing Punctuality. *Journal of the ACM* 43(1):116–146, 1996.
- [BFRS11] M. M. Bersani, A. Frigeri, M. Rossi, P. San Pietro. Completeness of the Bounded Satisfiability Problem for Constraint LTL. In *Reachability Problems*. LNCS 6945, pp. 58–71. 2011.
- [BK08] C. Baier, J.-P. Katoen. *Principles of Model Checking*. MIT Press, 2008.
- [BRS] M. M. Bersani, M. Rossi, P. San Pietro. A Tool for Deciding the Satisfiability of Continuous-time Metric Temporal Logic. To appear at TIME 2013.
- [DD07] S. Demri, D. D’Souza. An automata-theoretic approach to constraint LTL. *Inf. Comput.* 205(3):380–415, 2007.
- [FMMR12] C. A. Furia, D. Mandrioli, A. Morzenti, M. Rossi. *Modeling Time in Computing*. EATCS Monographs in Theoretical Computer Science. Springer, 2012.
- [HR99] Y. Hirshfeld, A. Rabinovich. Quantitative Temporal Logic. In *Computer Science Logic*. LNCS 1683, pp. 172–187. 1999.
- [HR05] Y. Hirshfeld, A. Rabinovich. Timer formulas and decidable metric temporal logic. *Information and Computation* 198(2):148 – 178, 2005.
- [Mic] Microsoft Research. Z3: An Efficient SMT Solver. <http://research.microsoft.com/en-us/um/redmond/projects/z3/>.
- [MNP06] O. Maler, D. Nickovic, A. Pnueli. From MITL to Timed Automata. In *Proc. of FORMATS*. LNCS 4202, pp. 274–289. 2006.
- [MP94] A. Morzenti, P. S. Pietro. Object-Oriented Logical Specification of Time-Critical Systems. *ACM TOSEM* 3(1):56–98, 1994.
- [PMS12] M. Pradella, A. Morzenti, P. San Pietro. Bounded Satisfiability Checking of Metric Temporal Logic Specifications. *ACM TOSEM*, 2012. To appear.
- [qtl] qtlsolver. available from qtlsolver.googlecode.com.
- [SRH02] P.-Y. Schobbens, J.-F. Raskin, T. A. Henzinger. Axioms for real-time logics. *Theor. Comput. Sci.* 274(1-2):151–182, 2002.