

Satisfiability Checking of Multi-Variable TPTL with Unilateral Intervals Is PSPACE-Complete

Shankara Narayanan Krishna ✉ 

IIT Bombay, Mumbai, India

Khushraj Nanik Madnani ✉ 

MPI-SWS, Kaiserslautern, Germany

Rupak Majumdar ✉ 

MPI-SWS, Kaiserslautern, Germany

Paritosh Pandya ✉ 

IIT Bombay, Mumbai, India

Abstract

We investigate the decidability of the $0, \infty$ fragment of Timed Propositional Temporal Logic (TPTL). We show that the satisfiability checking of $\text{TPTL}^{0, \infty}$ is PSPACE-complete. Moreover, even its 1-variable fragment ($1\text{-TPTL}^{0, \infty}$) is strictly more expressive than Metric Interval Temporal Logic (MITL) for which satisfiability checking is EXPSpace complete. Hence, we have a strictly more expressive logic with computationally easier satisfiability checking. To the best of our knowledge, $\text{TPTL}^{0, \infty}$ is the first multi-variable fragment of TPTL for which satisfiability checking is decidable without imposing any bounds/restrictions on the timed words (e.g. bounded variability, bounded time, etc.). The membership in PSPACE is obtained by a reduction to the emptiness checking problem for a new “non-punctual” subclass of Alternating Timed Automata with multiple clocks called Unilateral Very Weak Alternating Timed Automata ($\text{VWATA}^{0, \infty}$) which we prove to be in PSPACE. We show this by constructing a simulation equivalent non-deterministic timed automata whose number of clocks is polynomial in the size of the given $\text{VWATA}^{0, \infty}$.

2012 ACM Subject Classification Theory of computation → Logic

Keywords and phrases TPTL, Satisfiability, Non-Punctuality, Decidability, Expressiveness, ATA

Digital Object Identifier 10.4230/LIPIcs.CONCUR.2023.23

Related Version *Full Version*: <https://arxiv.org/abs/2309.00386>

Funding This research was sponsored in part by the Deutsche Forschungsgemeinschaft project 389792660 TRR 248—CPEC.

Acknowledgements We thank Tom Henzinger for encouraging us to explore non-punctual subclasses for multi-clock TPTL and ATA. We also thank Hsi-Ming Ho for an insightful discussion.

1 Introduction

Metric Temporal Logic ($\text{MTL}[\mathbf{U}_I, \mathbf{S}_I]$) and Timed Propositional Temporal Logic ($\text{TPTL}[\mathbf{U}_I, \mathbf{S}_I]$) are natural extensions of Linear Temporal Logic (LTL) for specifying real-time properties [6]. MTL extends the U and S modality of LTL by associating a time interval with these. Intuitively, $a\mathbf{U}_I b$ is true at a point in the given behaviour iff event a keeps on occurring until at some future time point within relative time interval I , event b occurs. (Similarly, $a\mathbf{S}_I b$ is its mirror image specifying the past behaviour.) On the other hand, TPTL uses freeze quantifiers to store the current time stamp. A freeze quantifier [4, 6] has the form $x.\varphi$ with freeze variable x (also called a clock [7, 27]). When it is evaluated at a point i on a timed word, the time stamp of i (say τ_i) is frozen or registered in x , and the formula φ is evaluated using this value for x . Variable x is used in φ in a constraint of the form $T - x \in I$; this constraint, when evaluated at a point j , checks if $\tau_j - \tau_i \in I$, where τ_j is the time stamp at point j . Here T can be seen as a special variable giving the timestamp of the present



© Shankara Narayanan Krishna, Khushraj Nanik Madnani, Rupak Majumdar, and Paritosh K. Pandya; licensed under Creative Commons License CC-BY 4.0

34th International Conference on Concurrency Theory (CONCUR 2023).

Editors: Guillermo A. Pérez and Jean-François Raskin; Article No. 23; pp. 23:1–23:18



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

point. For example, the formula $\varphi = Fx.(a \wedge F(b \wedge T - x \in [1, 2] \wedge F(c \wedge T - x \in [1, 2])))$ asserts that there is a point i in the future where a holds and in its future there is a b within interval $[1, 2]$ followed by a c within interval $[1, 2]$ from i . In this paper, we restrict ourselves to future time modalities only. Hence, we use the term MTL and TPTL for MTL[U_I] and TPTL[U], respectively, and MTL+Past and TPTL+Past for MTL[U_I, S_I] and TPTL[U, S], respectively. We also confine ourselves to the *pointwise* interpretation of these logics [7].

While these logics are natural formalisms to express real-time properties, it is unfortunate that both the logics have an undecidable satisfiability checking problem, making automated analysis of these logics difficult in general. Exploring natural decidable variants of these logics has been an active area of research since their advent [5, 31, 13, 35, 30, 14, 15]. One of the most celebrated such logics is the *Metric Interval Temporal Logic* (MITL) [1], a subclass of MTL where the timing intervals are restricted to be non-punctual i.e. non-singular (intervals of the form $\langle x, y \rangle$ where $x < y$). The satisfiability checking for MITL formulae is EXPSpace complete [1] (the result also holds for MITL + Past).

Every formula in MTL can be expressed in the 1-variable fragment of TPTL (denoted 1-TPTL). Moreover, the above-mentioned property φ is not expressible in MTL + Past [26]. Hence, 1-TPTL is strictly more expressive than MTL [27, 7]. The Logic 1-TPTL can also express MTL augmented with richer counting and Pnueli modalities. Hence, TPTL is a logic with high expressive power. However, decidable fragments of TPTL are harder to find. While 1-TPTL has decidable satisfiability over finite timed words [10] (albeit with non-primitive recursive complexity), it is undecidable over infinite words [25]. There are no known fragments of multi-variable TPTL which are decidable (without artificially restricting the class of timed words). In this paper, we propose one such logic, which is efficiently decidable over both finite and infinite timed words.

We propose a fragment of TPTL, called TPTL_{0,∞}, where, for any formula ϕ in negation normal form, each of its closed subformula κ has *unilateral* intervals; that is, intervals of the form $\langle 0, u \rangle$, or of the form $\langle l, \infty \rangle$ (where $\langle \in \{[, ($ and $\rangle \in \{],)\}$). The main result of this paper is to show that satisfiability checking for TPTL_{0,∞} is PSPACE complete. Moreover, we show that even the 1-variable fragment of this logic is strictly more expressive than MITL. PSPACE completeness for satisfiability checking is proved as follows: We define a sub-class of Alternating Timed Automata (ATA [24] [21]) called *Very Weak Alternating Timed Automata with Unilateral Intervals* (VWATA^{0,∞}), and show that VWATA^{0,∞} have PSPACE -complete emptiness checking. A language preserving reduction from TPTL_{0,∞} to VWATA^{0,∞}, similar to [10, 24, 34], completes the proof. To our knowledge, VWATA^{0,∞} is amongst the first known fragment of multi-clock alternating timed automata (ATA) with efficiently decidable emptiness checking. Thus, we believe that TPTL_{0,∞} and VWATA^{0,∞} are interesting novel additions to logics and automata for real-time behaviours.

One of the key challenges in establishing the decidability of VWATA^{0,∞} is to show that the configuration sizes can be bounded. In an ATA, a configuration can be unboundedly large owing to several conjunctive transitions, each spawning a state with a new clock valuation. We provide a framework for compressing the configuration sizes of VWATA^{0,∞} based on simulation relations amongst states of the VWATA^{0,∞}. We then prove that such compression yields a simulation-equivalent transition system whose configuration sizes are bounded. This bound allows us to give a subset-like construction resulting in a simulation equivalent (hence, language equivalent) timed automata with polynomially many clocks.

The paper is organized as follows. Section 2 defines the TPTL and ATA, and $(0, \infty)$ fragments of these formalisms. In Section 3, we prove the PSPACE emptiness checking of TPTL^{0,∞}. Section 4 discusses the expressiveness of TPTL_{0,∞}. Section 5 concludes our work with a discussion on the implication of our work in the field of timed logics and some interesting problems that we leave open.

2 Preliminaries

Let $\mathbb{Z}, \mathbb{Z}_{\geq 0}, \mathbb{N}, \mathbb{R}, \mathbb{R}_{\geq 0}$ respectively denote the set of integers, non-negative integers, natural numbers (excluding 0), real numbers, and non-negative real numbers. Given a sequence $\mathbf{a} = a_1 a_2 \dots$, $\mathbf{a}[i] = a_i$ denotes the i^{th} element of the sequence, $a[i..j]$ represents $a_i a_{i+1} \dots a_j$, $a[i..]$ represents $a_i a_{i+1} \dots$ and $a[..i]$ represents $a_1 a_2 \dots a_i$. Let \mathcal{I}_{int} be the set of all the open, half-open, or closed intervals (i.e. convex subsets of real numbers), such that the endpoints of these intervals are in $\mathbb{N} \cup \{0, \infty\}$. Intervals of the form $[x, x]$ are called punctual; a non-punctual interval is one which is not punctual. For, $\langle \in\{(\cdot, [\cdot]) \text{ and } \cdot\in\{(\cdot, \cdot)\} \rangle$, an interval of the form $\langle 0, u \rangle$ for $u > 0$ is called *right-sided* while an interval of the form $\langle l, \infty \rangle$ is called *left-sided*. A *unilateral* interval is either left-sided or right-sided. Let $\mathcal{I}_{\text{int}}^0, \mathcal{I}_{\text{int}}^\infty \subseteq \mathcal{I}_{\text{int}}$ respectively be the set of all *right sided* and *left sided* intervals of the form $\langle 0, u \rangle, \langle l, \infty \rangle$, for any $l, u \in \mathbb{Z}_{\geq 0}$. Let $\mathcal{I}_{\text{int}}^{0,\infty} = \mathcal{I}_{\text{int}}^0 \cup \mathcal{I}_{\text{int}}^\infty$. For $\tau \in \mathbb{R}$ and interval $\langle a, b \rangle$, $\tau + \langle a, b \rangle$ stands for the interval $\langle \tau + a, \tau + b \rangle$.

Timed Words. Let Σ be a finite alphabet. A finite (infinite) word over Σ is a finite (infinite) sequence over Σ . The set of all the finite (infinite) words over Σ is denoted by Σ^* (Σ^ω). A finite timed word ρ over Σ is a finite sequence of pairs $(\sigma, \tau) \in (\Sigma \times \mathbb{R}_{\geq 0})^*$: $\rho = (\sigma_1, \tau_1), \dots, (\sigma_n, \tau_n)$ where $\tau_i \leq \tau_j$ for all $1 \leq i \leq j \leq n$. Let $\text{dom}(\rho) = \{1, 2, \dots, n\}$ be the set of points in ρ . Likewise, an infinite timed word is an infinite sequence $\rho = (\sigma_1, \tau_1)(\sigma_2, \tau_2) \dots \in (\Sigma \times \mathbb{R}_{\geq 0})^\omega$, where $\sigma_1 \sigma_2 \dots \in \Sigma^\omega$, and $\tau_1 \tau_2 \dots$ is a monotonically increasing infinite sequence of real numbers approaching ∞ (i.e. non-zeno). A finite (infinite) timed language is a set of all finite (infinite) timed words over Σ denoted $T\Sigma^*$ ($T\Sigma^\omega$).

Timed Propositional Temporal Logic (TPTL). The logic TPTL extends LTL with freeze quantifiers and is evaluated on timed words. Formulae of TPTL are built from a finite alphabet Σ using Boolean connectives, as well as the temporal modalities of LTL. In addition, TPTL uses a finite set of real-valued variables called freeze variables or *clocks* $X = \{x_1, \dots, x_n\}$. Let $\nu : X \rightarrow \mathbb{R}_{\geq 0}$ represent a valuation assigning a non-negative real value to each clock. Without loss of generality, we work with TPTL in the negation normal form, where all the negations appear only with atomic formulae. Formulae of TPTL are defined as follows.

$$\varphi ::= a \mid \neg a \mid \top \mid \perp \mid x.\varphi \mid T - x \in I \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \varphi \text{U} \varphi \mid \text{G}\varphi$$

where $x \in X$, $a \in \Sigma$, $I \in \mathcal{I}_{\text{int}}$. T denotes the time stamp of the position where the formula is evaluated. The construct $x.\varphi$ is called a *freeze quantifier*, which stores in x , the time stamp of the current position and then evaluates φ . $T - x \in I$ is a constraint on the clock variable x , which checks if the time elapsed since the time x was frozen is in the interval I . Duals of Until; “Unless” and “Release” operators can be expressed using a G and an U operator without compromising on succinctness. **Notice** that, in aid of brevity, we will typically **abbreviate** subformula $T - x \in I$ to $x \in I$. For a timed word $\rho = (\sigma_1, \tau_1) \dots (\sigma_n, \tau_n)$, $i \in \text{dom}(\rho)$ and a TPTL formula φ , we define the satisfiability $\rho, i, \nu \models \varphi$ at a position i of ρ , given a valuation ν of the clock variables.

$$\begin{array}{lll} \rho, i, \nu \models a & \iff & \sigma_i = a, \\ \rho, i, \nu \models x.\varphi & \iff & \rho, i, \nu[x \leftarrow \tau_i] \models \varphi, \\ \rho, i, \nu \models T - x \in I & \iff & \tau_i - \nu(x) \in I, \\ \rho, i, \nu \models \text{G}\varphi & \iff & \forall j > i, \rho, j, \nu \models \varphi, \\ \rho, i, \nu \models \varphi_1 \text{U} \varphi_2 & \iff & \exists j > i, \rho, j \models \varphi_2, \text{ and } \forall i < k < j, \rho, k \models \varphi_1. \end{array}$$

The F and Next operator is defined in terms of U; $F\phi = \top U\phi$ and $\text{Next}\phi = \perp U\phi$. $\mathbf{0}$ denotes a valuation that maps every variable to 0. A TPTL formula φ is said to be closed iff every variable x used in the timing constraint is quantified (or bound) by a freeze quantifier. A formula that is not closed is *open*. Similarly, in any formula φ , a constraint of the form $x \in I$ is open if x is not quantified. For example, $x.y.(aU(b \wedge x \in (1, 2) \wedge y \in (2, 3)))$ is a closed formula while $x.(a \wedge y \in (2, 3))Uy.(b \wedge x \in (1, 2))$ is open as the clock y used in the underlined clock constraint is not in the scope of a freeze quantifier for y . Moreover, the underlined constraint $y \in (2, 3)$ is an open constraint. Notice that open constraints appear only (and necessarily) in open formulae. Satisfaction of closed formulae is independent of the clock valuation; that is, if ψ is a closed formula, then for a timed word ρ and a position i in ρ , either for every valuation ν , $\rho, i, \nu \models \psi$; or for every valuation ν , $\rho, i, \nu \not\models \psi$. Hence, for a closed formula ψ , we drop the valuation ν while evaluating satisfaction, and simply write $\rho, i \models \psi$. As an example, the closed formula $\varphi = x.(aU(bU(c \wedge x \in [1, 2])))$ is satisfied by the timed word $\rho = (a, 0)(a, 0.2)(b, 1.1)(b, 1.9)(c, 1.91)(c, 2.1)$ since $\rho, 1 \models \varphi$. The word $\rho' = (a, 0)(a, 0.3)(b, 1.4)(c, 2.1)(c, 2.5)$ does not satisfy φ . However, $\rho', 2 \models \varphi$: if we start from the second position of ρ' , the value 0.3 is stored in x by the freeze quantifier, and when we reach the position 4 of ρ' with $\tau_4 = 2.1$ we obtain $T - x = 2.1 - 0.3 \in [1, 2]$.

Given any closed TPTL formula φ , its language, $L(\varphi) = \{\rho \mid \rho, 1 \models \varphi\}$, is set of all the timed words satisfying it. We say that a closed formula φ is satisfiable iff $L(\varphi) \neq \emptyset$.

Size of a TPTL formula. Given a TPTL formula φ , the size of φ denoted by $|\varphi|$ is defined as $B + M + C$ where B is the number of Boolean operators in φ , M is the number of temporal modalities (G, U, Next, F) and freeze quantifiers in φ , and C is obtained by multiplying the number of time constraints in φ with $2 \times (\lceil \log(c_{max}) \rceil + 1)$ where c_{max} is the maximal constant appearing in the time constraints of φ . For example, for $\varphi = x.(a \wedge bU(c \vee x \leq (1, 2)))$, $|\varphi| = 2 + 2 + 2 \times (1 + 1) = 8$ as it contains two boolean operators, one temporal modality, one freeze quantifier and one timing constraint where $c_{max} = 2$.

The subclass of TPTL that uses **only k-clock variables** is known as **k-TPTL**. By [10] [25], satisfiability checking for 1-TPTL is decidable over finite models but non-primitive recursive hard, and undecidable over infinite models. Satisfiability checking for 2-TPTL is undecidable over both finite and infinite models [6] [18]. Towards the main contribution of this paper, we propose a “non-punctual” fragment of TPTL with unilateral intervals, called TPTL^{0,∞}, and show that its satisfiability checking is decidable with multiple variables over both finite and infinite timed words (PSPACE -complete). Further, 1-TPTL^{0,∞} is already more expressive than MITL, which has an EXPSpace -complete satisfiability checking.

2.1 Multi-clock TPTL with unilateral intervals: TPTL^{0,∞}

We say that a formula φ is of the type \leq (\geq), iff all the intervals appearing in the open constraints of φ are in $\mathcal{I}_{\text{int}}^0$ ($\mathcal{I}_{\text{int}}^\infty$). Notice that a closed formula belongs to both types \leq and \geq . There are open formulae that are neither of type \leq nor \geq . A TPTL formula φ in negation normal form is a TPTL^{0,∞} formula iff every subformula of φ is either of the type \leq or \geq . For example, $x.y.(aU(bU(c \wedge x < 3 \wedge y \leq 2 \wedge x.(Next(c \wedge x > 1))))))$ is a TPTL^{0,∞} formula since there is no subformula that doesn't belong to either types \leq or \geq . However, $x.y.(aU(b \wedge x \leq 3 \wedge y \geq 5))$ is not TPTL^{0,∞}, since $(b \wedge x \leq 3 \wedge y \geq 5)$ is of neither type \leq or \geq as the open constraints within this subformula use both left-sided as well as right-sided intervals. This restriction is inspired by that of MITL^{0,∞}. Any MITL^{0,∞} formula can be expressed in 1-TPTL^{0,∞} by applying the same reduction from MITL to 1-TPTL (see Remark 15). Next, we introduce alternating timed automata which are useful in proving the main result, i.e., Theorem 2.

2.1.1 Alternating Timed Automata

An Alternating Timed Automata (ATA) is a 7-tuple $\mathcal{A} = (Q, \Sigma, \delta, q_0, Q_{\text{acc}}, X, \mathcal{G})$, where, Q is a finite set of locations, X is a finite set of clock variables, \mathcal{G} is a finite set of guards of the form $x \in I$ where $I \in \mathcal{I}_{\text{int}}$ and $x \in X$, δ is a transition function, $q_0 \in Q$ is the initial location, and $Q_{\text{acc}} \subseteq Q$ is a set of accepting locations. The transition function is defined as $\delta : Q \times \Sigma \mapsto \Phi(Q, \mathcal{G})$ where $\Phi(Q, \mathcal{G})$ is defined by the grammar $\varphi ::= \top | \perp | \varphi_1 \wedge \varphi_2 | \varphi_1 \vee \varphi_2 | q|x \in I | Y.q$ with $q \in Q$, $x \in X$, $(x \in I)$ is a guard in \mathcal{G} , $Y \subseteq X$, Y is not the empty set. \top, \perp respectively denote True and False. $Y.q$ is a *binding construct* which resets all clocks in Y to zero after taking the transition. Let $p, q \in Q$ and $Y \subseteq X$. We say that there is a transition from p to q iff q appears in $\delta(p, b)$ for some $b \in \Sigma$. We say that there is a **strong reset transition**, **non-reset transition**, and a **Y-reset transition** from location p to q iff for some $b \in \Sigma$, $X.q$, q , and $Y.q$, respectively, appears in $\delta(p, b)$ for some $b \in \Sigma$. The 1-clock restriction of ATA has been considered in [24] and [21].

Evaluation of $\Phi(Q, \mathcal{G})$. Given an ATA \mathcal{A} , a state s is defined as a pair consisting of a location and a valuation over X , i.e., $s \in Q \times \mathcal{V}_X$. A configuration C of an ATA is a finite set of states. Let S and \mathcal{C} respectively denote the set of all states and configurations of \mathcal{A} . A configuration C and a clock valuation ν define a Boolean valuation for $\Phi(Q, \mathcal{G})$ as follows:

$$\begin{array}{l|l} C \models_{\nu} q \text{ iff } (q, \nu) \in C, & C \models_{\nu} Y.q \text{ iff } (q, \nu) \in C, \text{ and } \forall x \in Y. \nu(x) = 0, \\ C \models_{\nu} x \in I \text{ iff } \nu(x) \in I, & C \models_{\nu} \varphi_1 \wedge \varphi_2 \text{ iff } C \models_{\nu} \varphi_1 \wedge C \models_{\nu} \varphi_2, \\ C \models_{\nu} \top \text{ for all } C \in \mathcal{C}, & C \models_{\nu} \varphi_1 \vee \varphi_2 \text{ iff } C \models_{\nu} \varphi_1 \vee C \models_{\nu} \varphi_2. \end{array}$$

Finally, $C \not\models_{\nu} \perp$ for all possible configurations. We say that C is a minimal model for $\varphi \in \Phi(Q, \mathcal{G})$ with respect to ν (denoted by $C \models_{\nu}^{\text{min}} \varphi$) iff $C \models_{\nu} \varphi$ and no proper subset C' of C is such that $C' \models_{\nu} \varphi$. See Figure 1 in the full version for the graphical representation of the ATA.

Semantics of ATA. Given a state $s = (q, \nu)$, a time delay $t \in \mathbb{R}_{\geq 0}$ and $a \in \Sigma$, the successors of $s = (q, \nu)$ on time delay t followed by a is any configuration C such that $C \models_{\nu+t}^{\text{min}} \delta(q, a)$. $\text{Succ}_{\mathcal{A}}^{st}(s, t, a)$ is the set of all such successors. The notion of a successor is extended to a configuration in a straightforward manner. A configuration C' is a successor of configuration $C = \{s_1, s_2, \dots, s_k\}$ on time delay t and $a \in \Sigma$ (denoted by $C \xrightarrow{(t, a)}_{\mathcal{A}} C'$) iff $C' = C_1 \cup \dots \cup C_k$ such that $\forall 1 \leq i \leq k, C_i \in \text{Succ}_{\mathcal{A}}^{st}(s_i, t, a)$. We denote by $\text{Succ}_{\delta}(C, t, a)$ set of all such successors C' .

The initial configuration is defined by $C_{\text{init}} = \{(q_0, \mathbf{0})\}$, and a configuration C is accepting iff for all $s \in C$, s is an accepting state, that is $s = (q, \nu)$ for $q \in Q_{\text{acc}}$. Let \mathcal{C}_{acc} be the set of all the accepting configurations. Hence, the empty configuration is an accepting configuration. We define the semantics of ATA using a Labelled Transition System (LTS). An LTS is a 5-tuple $T = (S, s_0, \Sigma, \delta, S_f)$, where S is a finite or infinite set of states, $s_0 \in S$ is the initial state, Σ is set of symbols, $\delta : S \times \Sigma \times S$ is a transition relation, and $S_f \subseteq S$ is a set of final states. A (finite) run \mathbf{R} of an LTS is a (finite) sequence of the form $s_0, a_1, s_1, a_2, s_2, a_3, \dots$ where $s_1, s_2, \dots \in S$ are states of T , and a_1, a_2, \dots are symbols in Σ such that for all $i > 0$, $s_i \in \delta(s_{i-1}, a_i)$. We say that a run $R = s_0, a_1, s_1, a_2, s_2, a_3, \dots$ visits a state s (or visits a set of states S') iff the sequence R contains s (or contains states in S'). A run is said to be accepting iff it ends in some state $s \in S_f$. Similarly, an infinite run is said to be Büchi accepting iff it visits S_f infinitely often.

Runs of $\mathcal{A} = (Q, \Sigma, \delta, q_0, \mathbf{Q}_{acc}, X, \mathcal{G})$ starting from a configuration C are the runs of LTS $TS(\mathcal{A}, C) = (\mathcal{C}, C, \mathbb{R}_{\geq 0} \times \Sigma, \rightarrow, \mathcal{C}_{acc})$. Notice that the states of LTS $TS(\mathcal{A}, C)$ are configurations of \mathcal{A} (i.e., a set of states of \mathcal{A} and not just the states of \mathcal{A}). Let $\rho = (a_1, \tau_1)(a_2, \tau_2) \dots$ be any timed word over Σ . We say that a run $R = C, (t_1, a_1), C_1, (t_2, a_2) \dots$ is produced by \mathcal{A} on ρ starting from a configuration C iff $C \xrightarrow{(t_1, a_1)} C_1 \xrightarrow{(t_2, a_2)} C_2 \dots$ where $t_i = \tau_i - \tau_{i-1}$ for $i > 0$ and $\tau_0 = 0$. Let $\mathcal{A}(\rho, C)$ be the set of all the runs produced by \mathcal{A} on ρ , starting from the configuration C . We denote $TS(\mathcal{A}, C_{init})$ as simply $TS(\mathcal{A})$. A run starting from the initial configuration C_{init} is called an initialized run. We denote $\mathcal{A}(\rho, C_{init})$ by $\mathcal{A}(\rho)$. ρ, i is said to be accepted (Büchi accepted) by \mathcal{A} starting with configuration C , denoted by $\rho, i \models \mathcal{A}, C$, iff there exists a run in $\mathcal{A}(\rho[i..], C)$ accepted (Büchi accepted) by $TS(\mathcal{A})$ (i.e., simulating \mathcal{A} on the suffix of ρ starting at position i we obtain an accepting run). We say that ρ is accepted by \mathcal{A} iff $\rho, 1 \models C_{init}$.

We define the finite (infinite) language of \mathcal{A} , denoted by $L_{fin}(\mathcal{A})$ ($L_{inf}(\mathcal{A})$), as a set of all the finite (infinite) timed words accepted by \mathcal{A} . When clear from context, we drop the subscript in L_{fin} and L_{inf} .

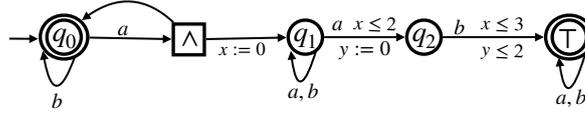
Non-Deterministic Timed Automata (NTA) is a subclass of ATA where $\Phi(Q, \mathcal{G})$ is restricted to be in disjunctive normal form (DNF), where each disjunct is of the form $(q \wedge x \in I)$ or $(X'.q \wedge x \in I)$. Hence, for any $s \in S$, $t \geq 0$, $a \in \Sigma$ and any configuration $C \in \text{Succ}_\delta^{st}(s, t, a)$ implies $C \leq 1$.

We call the ATA \mathcal{A} a **Very Weak ATA** (VWATA) iff (1) there is a partial order $\ll_{\mathcal{A}} \subseteq Q \times Q$ such that there is a transition from p to q iff $q \ll_{\mathcal{A}} p$, (2) all the self-loop transitions (transitions entering and exiting into the same location) are non-reset transitions, and (3) For every location q , there is at most one location $p \neq q$ such that there is a transition from p to q . Moreover, all the transitions from p to q reset the same set of clocks. This makes the transition diagram of VWATA a tree and not a DAG (excluding self-loops).

► **Remark 1.** In the literature, VWATA (also called Partially-Ordered Alternating Timed Automata in [20]) and their corresponding untimed version [9, 32] (also called as Linear [22], Linear-Weak [11], 1-Weak [28], and Self-Loop [33] Alternating Automata) are required to satisfy only conditions (1) and (2). It can be shown that condition (3) does not affect the expressiveness of the machine. We notice that this version of VWATA is enough to express TPTL formulae efficiently (linear in the size of TPTL formulae). In case of translation from TPTL to VWATA satisfying condition (3) the number of locations in the resulting ATA will depend on the size of the formula tree. On the other hand, the total number of locations depends on the formula DAG on similar translation from TPTL to VWATA satisfying only (1) and (2) making it exponentially more succinct. Hence, we consider a less succinct representation (i.e., tree or string, which is standard) of TPTL formulae for computing its size as compared to the DAG representation.

2.1.2 ATA with Unilateral Intervals: ATA^{0,∞}

Similar to the unilateral version of TPTL (i.e. TPTL^{0,∞}), we define a unilateral version of ATA, i.e., ATA^{0,∞} as follows. Let $\mathcal{A} = (Q, \Sigma, \delta, q_0, \mathbf{Q}_{acc}, X, \mathcal{G})$ be any ATA. Let \mathcal{G}^\geq (\mathcal{G}^\leq) be the subset of \mathcal{G} containing all the guards of the form $x \in I$ where $I \in \mathcal{I}_{int}^\infty$ ($I \in \mathcal{I}_{int}^0$). \mathcal{A} is said to be an ATA^{0,∞} iff, Q can be partitioned into Q^\geq and Q^\leq any transition exiting from any location $q \in Q^\geq$ ($q \in Q^\leq$) is guarded by a guard in \mathcal{G}^\geq (\mathcal{G}^\leq), and any transition from any location in Q^\geq to a location in Q^\leq , or vice-versa, is a strong reset transition. \mathcal{A} is said to be VWATA^{0,∞} iff it is an ATA^{0,∞}, and a VWATA. From this point onwards, for any set of locations Q of ATA^{0,∞}, Q^\geq and Q^\leq will denote partitions of Q satisfying the above condition.



■ **Figure 1** $VWATA^{0,\infty}$ equivalent to φ . Location q_i corresponds to the subformula φ_i : $\rho, i, \nu \models \varphi_i$ iff $\rho, i \models (q_i, \nu)$.

3 Satisfiability Checking for $TPTL^{0,\infty}$

This section is dedicated to proving the following main theorem of this paper.

► **Theorem 2.** *Satisfiability Checking for $TPTL^{0,\infty}$ is PSPACE -Complete*

PSPACE hardness follows from the hardness of satisfiability checking of the sublogics LTL and MITL^{0,∞} (see section 4.1 for the details on MITL^{0,∞}). To show membership in PSPACE we propose the following steps: (1) We reduce any given k - $TPTL^{0,\infty}$ formula φ , to an equivalent $VWATA^{0,\infty}$, \mathcal{A} , with k clock variables and at most $|\varphi| + 1$ number of locations. (2) We give a novel on-the-fly construction from any $VWATA^{0,\infty}$ to simulation equivalent NTA \mathbb{A} with exponential blow-up in the number of locations and polynomial blow-up in the number of clocks. Hence, the region automata corresponding to \mathbb{A} has at most exponentially many states, and thus each state can be represented in polynomial space.¹

► **Remark 3.** Notice that while the reduction from $VWATA^{0,\infty}$ to timed automata results in an exponential blow-up in the number of locations we can directly construct the region automaton of the corresponding timed automaton on-the-fly making sure that we need at most polynomial space to solve its emptiness checking problem.

We demonstrate our steps of construction using a running example. For the formal constructions please refer to the full version. In our running example, we start with the given formula $\varphi = G(\neg a \vee x.(F(a \wedge T - x \leq 2 \wedge y.Next(b \wedge T - x \leq 3 \wedge T - y \leq 2)))$.

3.1 $TPTL^{0,\infty}$ to $VWATA^{0,\infty}$

This step is a straightforward multi-clock generalization of translation from MTL and 1- $TPTL$ to 1-ATA in [24] and [10], respectively, (which are themselves timed generalization of reduction from LTL to Very Weak Alternating Automata [34] [9]). We give the reduction in the full version for completeness. The proof of equivalence is identical to that in [24] and [10] resulting in the following Theorem 4. We give the $VWATA^{0,\infty}$ corresponding to the formula φ of the running example in Figure 1. Hence, to prove the main theorem it suffices to show that emptiness checking for $VWATA^{0,\infty}$ is in PSPACE (i.e. Theorem 5).

► **Theorem 4.** *Any k variable $TPTL$ formula φ over Σ can be reduced to an equivalent $VWATA$, $\mathcal{A} = (Q, 2^\Sigma, \delta, init, Q_{acc}, X, \mathcal{G})$, with $|X| = k$, $|Q| \leq |\varphi| + 1$, and \mathcal{G} is the set of all the guards appearing in φ . Moreover, if φ is a $TPTL^{0,\infty}$ formula, then the \mathcal{A} is $VWATA^{0,\infty}$.*

¹ While one can argue about existence of a simple reduction from $TPTL^{0,\infty}$ to Recursive Memory Event Clock Logic of [17] using projections, we would still need to show that such a reduction requires only bounded memory which can be non-trivial, especially with multiple clocks. We believe that the automata-theoretic argument in this paper is a clean technique for proving such bounds.

3.2 Emptiness Checking for VWATA^{0,∞}

The following theorem is the main technical result.

► **Theorem 5.** *Emptiness Checking for VWATA^{0,∞} is in PSPACE .*

We give a translation from VWATA^{0,∞} $\mathcal{A} = (Q, \Sigma, \delta, q_0, Q_{\text{acc}}, X, \mathcal{G})$ to an equivalent timed automaton, $\mathbb{A} = (\mathcal{Q}, \Sigma, \Delta, q_0, Q_{\text{acc}}, \mathcal{X}, \mathcal{G})$, such that the transition system of \mathcal{A} (i.e., $TS(\mathcal{A})$) is simulation equivalent to that of \mathbb{A} (i.e., $TS(\mathbb{A})$). Hence, by the Proposition 6, $L(\mathcal{A}) = L(\mathbb{A})$.

Moreover, $Q = O(2^{\text{Poly}(Q)})$ and $|\mathcal{X}| = |X| \times |Q|$. Hence, the number of states in the corresponding region automaton is exponential to the size of \mathcal{A} (i.e. $O(2^{\text{Poly}(|Q|, |X|)}) \times (2 \times c_{\text{max}} + 1)$ where c_{max} is the maximum constant used in the constraints appearing in \mathcal{G}). Hence, each state of the region automata (when encoded in binary) can be represented in polynomial space proving membership in PSPACE . We prove the above by giving a translation from VWATA^{0,∞} to timed automata with polynomial blowup in the number of clocks and exponential blowup in the set of locations. As a side-effect, we also show that emptiness checking for 1-ATA^{0,∞} is in PSPACE (using the same construction) generalizing the result of [16]. We first briefly discuss the concept of simulation relations and preorders.

3.2.1 Simulation Relations and Preorder

We fix a pair of labeled transition system, $TS^1 = (S^1, s_0^1, \Sigma, \delta^1, S_f^1)$ and $TS^2 = (S^2, s_0^2, \Sigma, \delta^2, S_f^2)$. A relation $\preceq \subseteq S^1 \times S^2$ is a simulation relation iff (1) $s_0^1 \preceq s_0^2$, (2) for every $s_1 \preceq s_2$, (2.1) if $s_1 \in S_f^1$ then $s_2 \in S_f^2$, and (2.2) for every $a \in \Sigma$, for every $s'_1 \in \delta(s_1, a)$ there exists $s'_2 \in \delta(s_2, a)$ such that $s'_1 \preceq s'_2$. If $s_1 \preceq s_2$, then we say that s_2 simulates s_1 wrt \preceq .

Let $S = S^1 \cup S^2$. Notice that simulation relations are closed under union. Hence, there is a unique maximal simulation relation, $\leq \subseteq S \times S$, which is the union of all the simulation relations amongst states of TS^1 and TS^2 (i.e. all the simulation relations between TS^1 and itself, between TS^2 and itself, and from TS^1 to TS^2 and vice-versa). Notice that \leq is a preorder relation (i.e. reflexive and transitive), and hence also called simulation preorder. Similarly, simulation equivalence relation, \cong is defined as the largest symmetric subset of simulation preorder, \leq . I.e., $s \cong s'$ iff $s \leq s'$ and $s' \leq s$. Hence, it is clear that \cong is an equivalence relation. If $s \leq s'$ we say that s' simulates s . Recall that the states of $TS(\mathcal{A}, C)$ for any ATA \mathcal{A} and its configuration C are configurations of \mathcal{A} . Then,

► **Proposition 6.** *Let \mathcal{A} and \mathcal{A}' be any ATA, and s_0, s'_0 be their initial states, respectively. $TS(\mathcal{A}, \{s\}) \leq TS(\mathcal{A}', \{s'\})$ implies $L_{\text{fin}}(\mathcal{A}) \subseteq L_{\text{fin}}(\mathcal{A}')$ and $L_{\text{inf}}(\mathcal{A}) \subseteq L_{\text{inf}}(\mathcal{A}')$. Hence, $TS(\mathcal{A}, \{s\}) \cong TS(\mathcal{A}', \{s'\})$ implies $L_{\text{fin}}(\mathcal{A}) = L_{\text{fin}}(\mathcal{A}')$ and $L_{\text{inf}}(\mathcal{A}) = L_{\text{inf}}(\mathcal{A}')$*

We fix an ATA $\mathcal{A} = (Q, \Sigma, \delta, q_0, Q_{\text{acc}}, X, \mathcal{G})$. Let C and C' be arbitrary configurations of \mathcal{A} . Let $\leq_{\mathcal{A}}, \cong_{\mathcal{A}}$ be the simulation preorder and simulation equivalence amongst configurations of \mathcal{A} . That is, $C \leq_{\mathcal{A}} C'$ iff C' simulates C , and $C \cong_{\mathcal{A}} C'$ iff C is simulation equivalent to C' in $TS(\mathcal{A})$, the transition system corresponding to ATA \mathcal{A} . Then, by Proposition 6:

► **Remark 7.** For any configuration C and C' of \mathcal{A} , $C \leq_{\mathcal{A}} C'$ implies $L(\mathcal{A}, C) \subseteq L(\mathcal{A}, C')$ and $C \cong_{\mathcal{A}} C'$ implies $L(\mathcal{A}, C) = L(\mathcal{A}, C')$.

► **Remark 8.** $C \supseteq C'$ implies $C \leq_{\mathcal{A}} C'$. Hence, for any timed word ρ , if $\rho, i \models \mathcal{A}, C$ then $\rho, i \models \mathcal{A}, C'$. Intuitively, the additional states in C (which are not appearing in C') impose extra obligations in addition to that imposed by states common in both C and C' which makes reaching the accepting configuration (hence accepting a timed word) harder from C . For formal proof, please refer to the full version.

► **Remark 9.** If $D' \subseteq C$ and $D \leq_{\mathcal{A}} D'$, then $(C \setminus D') \cup D \leq_{\mathcal{A}} C$. In other words, we can replace the states in D' with that in D in any configuration C , and get a configuration that is simulated by C . Hence, $L(\mathcal{A}, (C \setminus D') \cup D) \subseteq L(\mathcal{A}, C)$.

Proof outline of Remark 9. First, show that for any configurations E_1, E_2 , and E if both E_1 and E_2 individually simulate E , then $(E_1 \cup E_2)$ simulates E . Second, substitute $E_1 = C \setminus D', E_2 = D'$, and $E = (C \setminus D') \cup D$. By Remark 8, E_1 and D individually simulate E . $E_2 = D'$ simulates D is given. Hence, E_2 simulates E by transitivity of preorders. Thus, $(E_1 \cup E_2)$ simulates E proving our remark. For full proof please refer to the full version. ◀

Both the above remarks imply the following Proposition. We abuse the notation by writing $\{s\} \leq_{\mathcal{A}} \{s'\}$ as $s \leq_{\mathcal{A}} s'$.

► **Proposition 10.** *If $s, s' \in C$ and $s \leq_{\mathcal{A}} s'$ then $C \setminus \{s'\} \cong_{\mathcal{A}} C$.*

Proof. Notice that $(C \setminus \{s'\}) \cup \{s\} = C \setminus \{s'\}$. Hence, by Remark 9, $(C \setminus \{s'\}) \leq_{\mathcal{A}} C$. By Remark 8, $C \leq_{\mathcal{A}} C \setminus \{s'\}$. Hence proved. ◀

We use the above Proposition 10 and Lemma 14 (which holds for $\text{VWATA}^{0,\infty}$ and $1\text{-ATA}^{0,\infty}$) to bound the cardinality of the configuration preserving simulation equivalence. This bound on the cardinality of configurations will imply that we need to remember only a bounded number of clock values to simulate these configurations. Hence, we use this bound on the cardinality of the configurations to bound the number of clock copies required while constructing the required timed automaton.

3.2.2 Bounding Cardinality of Configurations

Intuition

We now discuss the intuition for the decidability of $\text{VWATA}^{0,\infty}$. The main reason for the undecidability of ATA or VWATA is due to the unboundedness of the configuration size. That is, the cardinality of the configurations could depend on the length of the timed word prefix read so far. Hence, we need to keep track of an unbounded number of clocks. This happens, because we can reset a clock x in one branch and not reset x in another branch while taking transitions. This is a result of transitions containing clauses of the form $(X_i.q_i \wedge X_j.q_j)$ where $X_i \neq X_j$ and $X_i, X_j \subseteq X$. That is, we get two states in the successive configuration each resetting a different set of clocks. Hence, we need to remember multiple values for clock variables that are reset in one branch and not in another. In case of $\text{ATA}^{0,\infty}$, we observe the following:

- **Observation 1** – Let $q \in \mathbb{Q}^{\geq}$. Due to the nature of constraints, i.e. $x_i \in (l, \infty)$, if we have a pair of states $(q, \nu_1), (q, \nu_2)$ in a configuration C , such that $\nu_1 \leq \nu_2$ (i.e. $\forall x \in X. \nu_1(x) \leq \nu_2(x)$), then any timing constraint that is satisfied by ν_1 will also be satisfied by ν_2 . Hence, any transition that can be taken by (q, ν_1) can also be taken by (q, ν_2) . Moreover, after taking the same transition (time delay followed by event-based transition) both (q, ν_1) and (q, ν_2) get states of the form (q', ν'_1) and (q', ν'_2) , respectively, in their successor configurations, such that $\nu'_1 \leq \nu'_2$ if $q' \in \mathbb{Q}^{\geq}$ and $\nu'_1 = \nu'_2 = \mathbf{0}$ if $q' \in \mathbb{Q}^{\leq}$. Hence, by Proposition 10, we can delete (q, ν_2) from C preserving simulation equivalence (and hence the language). A similar argument applies for $q \in \mathbb{Q}^{\leq}$.
- **Observation 2** – In 1-ATA , for any pair of valuations ν_1, ν_2 , either $\nu_1 \leq \nu_2$ or $\nu_2 \leq \nu_1$. Hence, on applying the reduction using Proposition 10 (and discussed in the previous bullet, i.e., Observation 1), we will always get a configuration, where each location appears at most once. Hence, the configuration size is bounded by the number of locations.

23:10 Satisfiability Checking for TPTL^{0,∞} Is PSPACE-Complete

- Observation 3 – But this is not necessarily the case for multiple clocks. This is because there could be unboundedly many incomparable valuations. For example, for 2-clocks $X = \{x, y\}$, consider the following family of configurations parameterized by m , $C_m = \{(q, x = 0.1 + nk, y = 0.9 - nk) | n \in \{0, \dots, m-1\}\}$ and $k = 0.8/m$. $|C_m| = m$ and all the clock valuations are incomparable. Notice $C_8 = \{(q, x = 0.1, y = 0.9), (q, x = 0.2, y = 0.8) \dots (q, x = 0.9, y = 0.1)\}$. Hence, as the second main step we show that, if \mathcal{A} is a VWATA^{0,∞}, and if we conservatively keep on compressing the configurations as discussed in Observation 1 (using Proposition 10), we will have boundedly many incomparable clock valuations. To be precise, we will have at most one copy of each location in the configuration. This is shown in Lemma 14.

Bounding Lemma

In this section, we will use the intuition in Observation 1 for constructing a simulation equivalent transition system for a given 1-ATA^{0,∞} and VWATA^{0,∞} whose states are configurations of given ATA \mathcal{A} with bounded cardinality. For the 1-ATA^{0,∞}, the intuition in Observation 2 guarantees the case. For the multi-clock VWATA^{0,∞}, the issues discussed in Observation 3 must be resolved. This is resolved in Lemma 14, the main contribution of this section. In what follows, assume \mathcal{A} to be an ATA^{0,∞}. We define relation \preceq amongst states of \mathcal{A} . For $\sim \in \{\leq, \geq\}$, let \preceq be defined between states such that $s \preceq s'$ iff $s = (q, \nu)$, $s' = (q, \nu')$, and if $q \in Q^\sim$ then $\nu' \sim \nu$. By Observation 1 we have Proposition 11. The formal proof appears in the full version.

► **Proposition 11.** $s \preceq s'$ implies $s \leq_{\mathcal{A}} s'$.

Given any configuration C , we define $\text{Red}_{\preceq}(C)$ as a configuration C' obtained from C , by deleting all states $s' \in C'$ if there exists a state $s \in C'$, such that $s \neq s'$, and $s \preceq s'$. Intuitively, we delete some information from a configuration that is redundant in deciding whether a timed behaviour from that state is accepted or not.

Let C_0 be the initial configuration of \mathcal{A} . Let $TS(\mathcal{A}) = (\mathcal{C}, C_0, (\mathbb{R}_{\geq 0} \times \Sigma), \rightarrow_{\mathcal{A}})$ be the transition system corresponding to \mathcal{A} . We define $\text{T}_{red}(\mathcal{A})$ as a transition system $\text{T}_{red}(\mathcal{A}) = (\mathcal{C}, C'_0, (\mathbb{R}_{\geq 0} \times \Sigma), \rightarrow_{\mathcal{A}, red})$ such that $C'_0 = \text{Red}_{\preceq}(C_0)$ and for any $C, C', D, D' \in \mathcal{C}$, $a \in \Sigma$, and $t \in \mathbb{R}_{\geq 0}$, $C \xrightarrow{(t, a)}_{\mathcal{A}} C'$ iff $D \xrightarrow{(t, a)}_{\mathcal{A}, red} D'$, $D = \text{Red}_{\preceq}(C)$, and $D' = \text{Red}_{\preceq}(C')$. By Proposition 12, $TS(\mathcal{A})$ is simulation equivalent to $\text{T}_{red}(\mathcal{A})$. The following Proposition is implied by Proposition 10 and 11.

► **Proposition 12.** $C \cong_{\mathcal{A}} \text{Red}_{\preceq}(C)$. Hence, $TS(\mathcal{A})$ and $\text{T}_{red}(\mathcal{A})$ are simulation equivalent.

► **Remark 13.** Any run R' is a run of $\text{T}_{red}(\mathcal{A})$ iff $R' = \text{lmg}(R)$ for some run R of \mathcal{A} , where $\text{lmg}(R)$ is defined as follows. $R = C_0 \xrightarrow{(t_0, a_0)}_{\mathcal{A}} C_1 \xrightarrow{(t_1, a_1)}_{\mathcal{A}} C_2 \dots$, we define $\text{lmg}(R)$ as run $R' = C''_0 \xrightarrow{(t_0, a_0)} C''_1 \xrightarrow{(t_1, a_1)} C''_2 \dots$ where $C''_0 = C'_0 = \text{Red}_{\preceq}(C_0)$ and $\forall i \geq 0. C''_i \xrightarrow{(t_i, a_i)}_{\mathcal{A}} C''_{i+1}$ and $C''_i = \text{Red}_{\preceq} C'_i$.

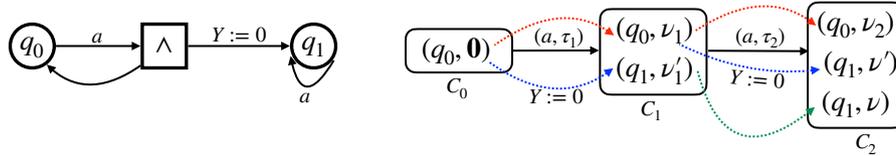
► **Lemma 14.** Let $\mathcal{A} = (Q, \Sigma, \delta, q_0, Q_{acc}, X, \mathcal{G})$ be either an 1-ATA^{0,∞} or VWATA^{0,∞}. Let R be a run of \mathcal{A} , and $R' = \text{lmg}(R) = C''_0(t_0, a_0)C''_1(t_1, a_1) \dots$, then for all $i \geq 1$, C''_i does not contain states (q, ν) and (q, ν') where $\nu \neq \nu'$ for any $q \in Q$. In other words, every location $q \in Q$ appears at most once in any configuration C''_i for any $i \geq 1$. Hence, $|C''_i| \leq |Q|$.

Proof (sketch). Notice that if \mathcal{A} was 1-ATA^{0,∞}, the above statement is straightforward as no two clock valuations are incomparable in the case of 1-clock. We now show the same for \mathcal{A} being a multi-clock VWATA^{0,∞}. We just present intuition behind the proof idea. A

formal proof is proved using DAG semantics of ATA and can be found in the full version. We prove this by contradiction. Assumption 1 - Suppose k is the smallest number such that C_k'' contains two copies of some location $q \in Q$. Hence, there exists ν and ν' such that ν' is incomparable to ν and $(q, \nu), (q, \nu') \in C_k''$. Then, the following cases are possible:

Case 1 - Both $(q, \nu), (q, \nu')$ appeared from the same location p in C_{k-1}'' . But, by condition (3) of VWATA, all the transitions from location p to location q reset the same set of clocks. Moreover, by assumption 1, location p appears at most once in C_{k-1}'' . Let $(p, \nu_p) \in C_{k-1}''$. Then both the clock valuations ν and ν' should be identical as they result from the same state (p, ν_p) resetting the same set of clocks.

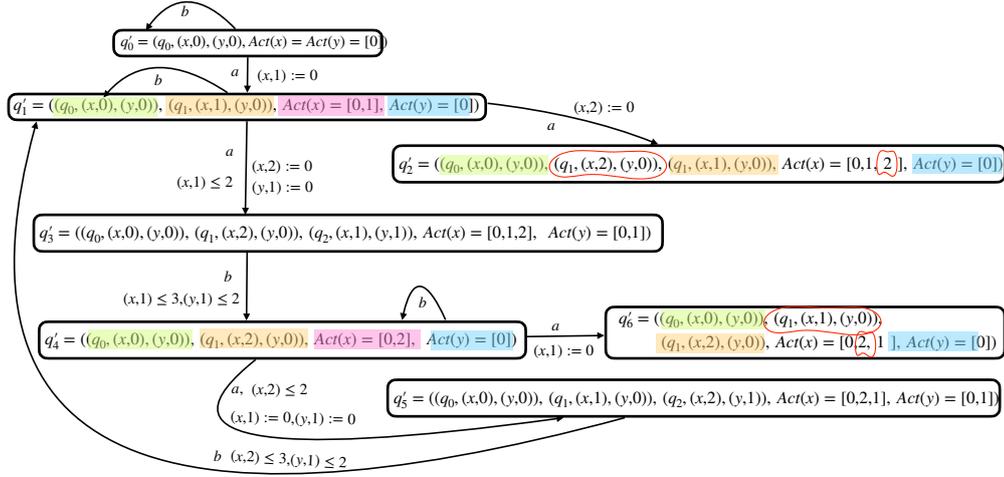
Case 2 - $(q, \nu), (q, \nu')$ appeared from distinct location (p, ν_{k-1}) and (p', ν'_{k-1}) in C_{k-1}'' . By condition (3) of VWATA there is at most one location $q' \neq q$ from which there are transitions entering location q . Moreover, all these transitions reset the same set of clocks. Hence, one of p and p' has to be q . Wlog $p = q$. It suffices to show that whenever such a case occurs, the clock valuation of the state that results from the self-loop (in this case ν) is always greater than or equal to the valuation from the other (in this case ν') (Statement 1). Hence, $\nu' \leq \nu$ which leads to a contradiction. We just present the intuition with an example. Let $\rho = (a_1, \tau_1), (a_2, \tau_2)$. Suppose, $(q_0, \mathbf{0})$ is the initial location of the automaton as drawn in Figure 2. Let $k = 2$. Notice the run in the Figure, $C_1 = \{(q_0, \nu_1), (q_1, \nu'_1)\}$ where if $x \in X'$, $\nu'_1(x) = 0 \leq \nu_1(x) = \tau_1$. Else, $\nu_1(x) = \nu'_1(x) = \tau_1$. Similarly, $C_2 = \{(q_0, \nu_2), (q_1, \nu'), (q_1, \nu)\}$, where (q_1, ν) results from the self loop and (q_1, ν') results from the transition from q_0 . Hence, if $x = X'$, $\nu(x) = 0 \leq \nu'(x) = \tau_2 - \tau_1$. Else, $\nu_1(x) = \nu'_1(x) = \tau_2$. In other words, while reaching both (q_1, ν) and (q_1, ν') from the initial configuration, the same set of clock X' was reset. But, in the case of the former, they were reset before the latter. Hence, ν and ν' agree on all the clock values not in X' and $\nu \geq \nu'$ for all the clocks in X . Applying this argument inductively we can prove Statement 1. We believe it is more intuitive to prove the result using the DAG semantics of ATA. Hence, the full proof can be found in the full version, where we introduce the semantics too. ◀



■ **Figure 2** The red and green transitions denote those without resets, and the blue ones with resets. Notice the paths from C_0 to C_2 . The Blue-Green and Red-Blue path reset the same set of clocks Y . But the former resets the clocks earlier (in the first step) as compared to the latter (in the second step). Hence in the former, clocks in Y get a chance to progress between C_1 and C_2 . Moreover, both the paths should agree on the value of clocks not in Y as they are not reset in both these paths. Hence, $\nu' \leq \nu$.

3.2.3 From $VWATA^{0,\infty}$ to Timed Automata

In this section, we propose an on-the-fly construction from $VWATA^{0,\infty}$ to Timed Automata. The termination relies on Lemma 14. The main idea is to bind the number of active clocks using Lemma 14. Given a $VWATA^{0,\infty}$ or $1-ATA^{0,\infty}$, $\mathcal{A} = (Q, \Sigma, \delta, q_0, Q_{acc}, X, \mathcal{G}, Q^{\geq}, Q^{\leq})$ we get a timed automaton $\mathcal{A} = (\mathcal{Q}, \Sigma, \Delta, q'_0, \mathcal{Q}_{acc}, X \times \{0, \dots, |Q| - 1\}, \mathcal{G})$ and at every step we reduce the size of the location $q \in \mathcal{Q}$ preserving simulation equivalence. Let V be set of



■ **Figure 3** Steps in the construction of \mathbb{A} corresponding to our running example. With the color coding in q'_1, q'_2 , it is easy to see that q'_2 is same as q'_1 on removing the circled entities in q'_2 . Same with q'_4 and q'_6 .

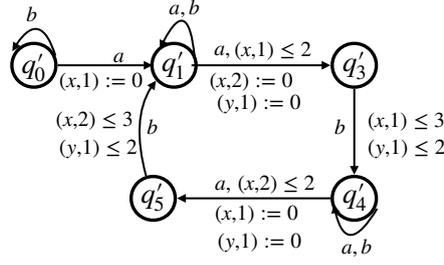
all the functions of the form $v : X \mapsto \{0, \dots, |Q| - 1\}$. Let \mathcal{L} be a set of all the functions from Q to $V \cup \{0\}$. Let **Active** be a set of all the functions from X to a sequence (without duplicate) over $\{0, 1 \dots |Q| - 1\}$. Then $\mathcal{Q} = \mathcal{L} \times \text{Active}$. Intuitively, we replace the bunch of conjunctive transitions C into a single transition, similar to the subset construction for converting Alternating Finite Automata (AFA) to Non-Deterministic Finite Automata (NFA). But notice that we can have clauses (or conjunctions) of the form $q \wedge X'.q'$. Hence, simple subset construction won't work as we need to spawn multiple copies of clocks in X' , wherein one of the elements of the new location $\{q, q'\}$ they are reset while in another they are not. In general, there could be an unbounded number of such clock copies required for a single clock, $x \in X$. But due to Lemma 14, if we make sure to compress the states (and hence remove redundant clocks), we need to keep at most $|Q|$ copies for each clock in X . In principle, we are constructing an NTA \mathbb{A} whose transition system $TS(\mathbb{A})$ is simulation equivalent to the LTS $T_{red}(\mathcal{A})$ (see the full version Proposition 19) and hence to input $\text{VWATA}^{0,\infty} TS(\mathcal{A})$. Thus, by Proposition 6, $L(\mathcal{A}) = L(\mathbb{A})$. We present the idea via our running example.

3.2.4 Construction on Running Example

Please refer to the $\text{VWATA}^{0,\infty}$ of our running example Figure 1. We now illustrate the construction on our running example. We start with location q_0 , with the 0^{th} copies of clock x and y . Hence

$$q'_0 = \{(q_0, (x, 0)(y, 0)), \text{Active}(x) = [0], \text{Active}(y) = [0]\}.$$

This corresponds to the configuration $C_0 = \{(q_0, \mathbf{0}_X)\}$ of \mathcal{A} . In the input automaton, the transitions from q_0 on a is defined by $\delta(q_0, a) = q_0 \wedge x.q_1$. Hence, we need to spawn a new copy of clock x as it is reset in one transition and not in another. We associate this new copy of clock x with the branch that resets x , i.e., this new clock x is associated with location q_1 . Hence, we have $\Delta(q'_0, a) = (x, 1).q'_1$ where $q'_1 = \{(q_0, (x, 0), (y, 0)), (q_1, (x, 1)(y, 0)), \text{Active}(x) = 0 \geq 1, \text{Active}(y) = 0\}$. Intuitively, q'_1 corresponds to the configurations of the form $C_1 = \{(q_0, \nu_0^1), (q_1, \nu_1^1)\}$ of \mathcal{A} , where $\nu_0^1(x) = \text{value of } (x, 0)$, $\nu_1^1(x) = \text{value of } (x, 1)$, and



■ **Figure 4** Final Automata after applying the reductions.

$\nu_0^1(y) = \nu_1^1(y) = \text{value of } (y, 0)$. We continue with this new location. Hence, we will consider the transitions from both q_0 and q_1 on a . The component $((q_0, (x, 0), (y, 0)))$ on a again spawns a new copy of clock x as it resets the clock in one while not resetting on self-loop, hence, getting $\{(q_0, (x, 0), (y, 0)), (q_1, (x, 2), (y, 0))\}$ (possibility 1 from q_0 , the only possibility). Notice that we spawned $(x, 2)$ as $(x, 1)$ is in use by q_1 already. The component $(q_1, (x, 1), (y, 0))$ will be computed using the transition function of input automaton, i.e. $\delta(q_1, a) = (y.q_2 \wedge x \leq 2) \vee q_1$. Here, we either stay at q_1 with the same set of clock copies as before (possibility 1 from q_1), or we need a new copy of y while simultaneously checking for the clock copy of x corresponding to location q_1 (i.e. $(x, 1)$) is ≤ 2 (possibility 2 from q_1). Combining the possibilities 1 from q_0 and q_1 we get, $\{(q_0, (x, 0), (y, 0)), (q_1, (x, 2), (y, 0)), (q_1, (x, 1), (y, 0)), \text{Active}(x)=[0 \geq 1 \geq 2], \text{Active}(y)=[0]\}$. But $q \in \mathcal{Q}^{\leq}$. Hence, if we can reach the accepting state from $(q_1, (x, 1), (y, 0))$ then we can reach the accepting state from $(q_1, (x, 2), (y, 0))$ too, as value of $(x, 1) \geq \text{value of } (x, 2)$ (this fact is also encoded in the $\text{Active}(x)$ sequence). Thus, $(q_1, (x, 2), (y, 0))$ can be removed from the new location without affecting simulation equivalence (and hence language equivalence). This corresponds to the removal of redundant states in the construction of the runs of $\mathbb{T}_{red}(\mathcal{A})$ from $T(\mathcal{A})$. Hence, after deletion we get $q'_2 = \{(q_0, (x, 0), (y, 0)), \overline{(q_1, (x, 2), (y, 0))}, (q_1, (x, 1), (y, 0)), \text{Active}(x)=[0 \geq 1 \geq 2], \text{Active}(y)=[0]\} = q'_1$.

Thus, combining result of the transition of q_0 on a and possibility 1 from q_1 we get $q'_1 = \{q_0, (x, 0), (y, 0), (q_1, (x, 1), (y, 0)), \text{Active}(x)=0 \geq 1, \text{Active}(y)=0\}$.

Combining results possibility 1 from q_0 and possibility 2 from q_1 , we get

$\{q_0, (x, 0), (y, 0), (q_1, (x, 2), (y, 0)), (q_2, (x, 1), (y, 1)) \text{Active}(x)=0 \geq 2 \geq 1, \text{Active}(y)=0 \geq 1.\} = q'_3$ if $(x, 1) \leq 2$. Note that each location from \mathcal{Q} appears at most once in q'_3 . Hence, there is no scope of reduction.

Combining the above two combination of possibilities, $\Delta(q'_1, a) = q'_1 \vee (y, 1).(q'_3 \wedge x \leq 2)$.

Continuing this we get the resulting NTA \mathbb{A} equivalent to the input formula ϕ . Notice that we are eliminating the conjunctive transitions using subset like construction and keeping the disjunctions as it is. Hence, after eliminating all the conjunctive transitions the reduced automata contains only disjunctions amongst different locations in the output formulae of the transitions giving an NTA. Refer to Figures 3, 4.

3.2.5 Worst Case Complexity

By construction in [2], the number of states in the region automata of $\mathbb{A} = W \leq |\mathcal{Q}| \times (|X| \times |\mathcal{Q}|)! \times 2 \times (c_{max} + 1)$ where c_{max} is the max constant used in the guards in \mathcal{G} and $|\mathcal{Q}| = |\mathcal{Q}| \times (|X|^{|\mathcal{Q}|} + 1) \times (|\mathcal{Q}|!)^{|X|}$. Hence, $W = O(2^{Poly(|A|)})$ implying that the emptiness could be checked in $\text{NPSpace} = \text{PSPACE}$. Notice that the state containing the location (L, Act) will only have to store the region information of active clocks, which, in practice, could be much less than the worst case. Hence, lazily spawning clock copies may result in NTA with much less number of clocks than the worst case (i.e. $|X| \times |\mathcal{Q}|$).

4 Expressiveness of TPTL^{0,∞}

We now compare the expressive power of 1-TPTL^{0,∞} with respect to that of MITL.

4.1 Metric Temporal Logic(MTL)

MTL is a real-time extension of LTL where the U modality is guarded with an interval. Syntax of MTL is defined as follows. $\varphi ::= a \mid \top \mid \varphi \wedge \varphi \mid \neg \varphi \mid \varphi \mathbf{U}_I \varphi$, where $a \in \Sigma$ and $I \in \mathcal{I}_{\text{int}}$. For a timed word $\rho = (\sigma_1, \tau_1)(\sigma_2, \tau_2) \dots (\sigma_n, \tau_n) \in T\Sigma^*$, a position $i \in \text{dom}(\rho)$, an MTL formula φ , the satisfaction of φ at a position i of ρ , denoted $\rho, i \models \varphi$, is defined as follows. We discuss only the semantics of temporal modalities. Boolean operators mean as usual. $\rho, i \models \varphi_1 \mathbf{U}_I \varphi_2$ iff $\exists j > i. \rho, j \models \varphi_2, \tau_j - \tau_i \in I$, and $\forall i < k < j. \rho, k \models \varphi_1$. As usual, $F_I(\phi) = \top \mathbf{U}_I \phi$, $G(\phi) = \neg F_I \neg \phi$, $\text{Next}_I \phi = \perp \mathbf{U}_I \phi$. The language of an MTL formula φ is defined as $L(\varphi) = \{\rho \mid \rho, 1 \models \varphi\}$. The subclass of MTL where the intervals I in the “until” modalities are restricted to be **non-punctual** is known as Metric Interval Temporal Logic (MITL). MITL^{0,∞} [1, 3, 13] is the subclass of MTL where intervals are restricted in $\mathcal{I}_{\text{int}}^{0,\infty}$. Satisfiability Checking for MITL (MITL^{0,∞}) is EXPSPACE -complete (PSPACE -complete) [4, 1, 3]. MITL is strictly more expressive than MITL^{0,∞} in pointwise semantics [12].

► **Remark 15.** Any MTL formula can be translated to an equivalent 1-TPTL (closed) formula using the following equivalence recursively. $\varphi_1 \mathbf{U}_I \varphi_2 \equiv x.(\varphi_1 \mathbf{U} \varphi_2 \wedge x \in I)$.

4.2 Expressiveness of TPTL^{0,∞}

► **Theorem 16.** 1-TPTL^{0,∞} is strictly more expressive than MITL.

Proof. Both MITL and 1-TPTL^{0,∞} are closed under all boolean operations. Hence, we just need to show that any formula of the form $\varphi' \mathbf{U}_I \varphi$ is expressible in 1-TPTL^{0,∞}. Notice that, any MITL formula $\varphi' \mathbf{U}_{[l,u]} \varphi \equiv [G_{[0,l]} \{\varphi' \wedge (\varphi' \mathbf{U} \varphi)\}] \wedge [F_{[l,l+1]} \varphi \vee F_{[l+1,l+2]} \varphi \dots F_{[u-1,u]} \varphi]$. (similar reduction applies for other kinds of intervals). $G_{[0,l]}(\varphi' \wedge (\varphi' \mathbf{U} \varphi))$ is already in MITL^{0,∞} (and hence in 1-TPTL^{0,∞} by remark 15). Hence, it suffices to encode modalities of the form $F_{[l,l+1]}$ using 1-TPTL^{0,∞} formula. Let $\rho = (a_1, \tau_1), (a_2, \tau_2) \dots$ be any timed word. Let $i \in \text{dom}(\rho)$ be any point. $\rho, i \models F_{[l,l+1]}(\varphi)$ iff there exists a point $i' > i$ such that $\tau_{i'} - \tau_i \in [l, l+1)$ and $\rho, i' \models \varphi$. ρ has a point i' within $[l, l+1)$ interval from i where φ holds iff there exist earliest such point j ($j \leq i'$) within $[l, l+1)$ from i where φ holds iff there is a point $j' > i$ such that $\tau_{j'} - \tau_i \geq l$ (i.e. $\rho, i \models \phi_0 = F_{[l,\infty)}(\varphi)$), and let j be the first point such that $\tau_j - \tau_i \geq l$, and $\rho, j \models \varphi$. Such a point exists due to ϕ_0 . Then:

- Case 1: Either there is no point strictly between i and j where φ holds. Then occurrence of j within $l+1$ can be expressed using formula, $\phi_1 = \neg F_{[0,l]} \varphi \wedge F_{[0,l+1]} \varphi$.
- Case 2: Or there exists a point k such that $\tau_j - \tau_k < 1$, $\tau_k - \tau_i \in [l-1, l)$, and $\rho, k \models \varphi$. Equivalently, i satisfies $\phi_2 = G_{[l-1,l]}(F_{[0,1]}(\varphi))$,
- Case 3: Or there exists a point k with $i < k < j$ such that $\tau_j - \tau_k \geq 1$, $\rho, k \models \varphi$, and $\forall k < k' < j. \rho, k' \not\models \varphi$.

(1) Such a point k satisfies $\phi_{\text{approach}} = \varphi \wedge G_{[0,1]}(\neg \varphi)$. Indeed a key property is that k , the last point in $[0, l)$ satisfying ϕ , satisfies ϕ_{approach} . By the definition of k , i.e., there are no occurrences of φ after k in $[0, l)$.

(2) Notice that any two point k_1 and k_2 satisfying ϕ_{approach} are at least a unit time apart. Hence, there could be at most l points satisfying ϕ_{approach} within $[0, l)$. Then, the following 1-TPTL^{0,∞} formula $\text{Count}_{\phi_{\text{approach}}}(n)$ with parameter n states that there are exactly n points, $1 \leq n \leq l$ within $[0, l)$ of point i where ϕ_{approach} holds. Here,

$$\text{Count}_{\phi_{\text{approach}}}(n) = \phi_{\geq n} \wedge \neg \phi_{\geq n+1}, \text{ where } \phi_{\geq n} = x.((\neg \phi_{\text{approach}}) \mathbf{U}(\phi_{\text{approach}} \wedge (\neg \phi_{\text{approach}}) \mathbf{U}(\phi_{\text{approach}} \wedge x < l) \dots)))$$

$\underbrace{\dots}_{n-3}$

Observe that for a given timed word and interval $[0, l]$ from i , there is a unique n satisfying this formula $\text{Count}\phi_{\text{approach}}(n)$.

(3) Using this n , the formula $\gamma(n, \varphi) = x.(\neg\phi_{\text{approach}}\mathbf{U}(\phi_{\text{approach}} \wedge \neg\phi_{\text{approach}}\mathbf{U}(\phi_{\text{approach}} \wedge \underbrace{\dots}_{n-3} \wedge ((\neg\phi_{\text{approach}})\mathbf{U}(\phi_{\text{approach}} \wedge \mathbf{G}(x \leq l \vee \neg\varphi) \wedge \mathbf{F}(\varphi \wedge x < l + 1)) \dots)))$ holds if after n occurrences of ϕ_{approach} (which gives point k), the next occurrence of φ occurs before time $l + 1$.

Hence, case 3 is characterized by the formula $\phi_3 = \bigvee_{n=1}^l \text{Count}\phi_{\text{approach}}(n) \wedge \gamma(n, \varphi)$.

Hence, the required formula $\psi = \phi_0 \wedge (\phi_1 \vee \phi_2 \vee \phi_3)$.

For strict containment of MITL, consider the formula $\beta = x.\mathbf{F}(b \wedge \mathbf{F}(b \wedge x \leq 1))$. This specifies, there exist at least two points within the next unit interval where b holds. [15, 29, 23] show that this formula is not expressible even in MTL. \blacktriangleleft

5 Discussion and Conclusion

Ferrère [8] proposed an extension of LTL with Metric Interval Regular Expressions called Metric Interval Dynamic Logic (MIDL) and showed it to be more expressive than EMITL of [35]. We claim that our proof of PSPACE completeness for 1-ATA^{0,∞} emptiness implies the same for MIDL_{0,∞} satisfiability strictly generalizing the results and techniques of [16] which proved the same for EMITL_{0,∞}. This resolves one of the “future directions” of [16].

Authors in [19] generalized the notion of non-punctuality to non-adjacency for 1-TPTL. We remark that unfortunately, this notion doesn’t help in making 2-TPTL decidable. Notice that $\varphi = \mathbf{G}x.\{\neg\phi \vee \mathbf{F}y.(\top \wedge x \in [1, 2] \wedge \mathbf{F}(\phi_1 \wedge x \in [1, 2] \wedge y \in [1, 2]))\} \equiv \mathbf{G}[\phi \rightarrow \mathbf{F}_{[1,1]}(\mathbf{F}_{[1,1]}\phi_1)]$. Because, for any point i where φ holds there is a point j in the future such that $\tau_j - \tau_i \in [1, 2]$, and from that point j there is a point k in the future where ϕ_2 holds such that $\tau_k - \tau_j \in [1, 2]$ and $\tau_k - \tau_i \in [1, 2]$. Solving the inequalities we get, $\tau_j - \tau_i = 1$ and $\tau_k - \tau_i = 2$. Hence, φ can express some restricted form of punctual timing properties which leads to the undecidability of satisfiability using encoding similar to [25].

MITL^{0,∞} was extended with Counting (TLC) and Pnueli (TLP) modalities by [15] to increase the expressiveness, meanwhile maintaining the decidability in EXPSpace and PSPACE, respectively. These logics TLP and TLC have the same expressive power in continuous semantics. While these logics were strictly more expressive than MITL in continuous semantics, in pointwise semantics they are incomparable. This is due to inexpressivity of arbitrary non-punctual metric constraints using unilateral metric interval constraints in pointwise semantics (see [16]). However, TLP and TLC properties are trivially expressible in TPTL^{0,∞} (one clock and nested until), making our logic strictly more expressive than these. As one of our future works, we would like to show that TLCI and TLPI (extensions of TLP and TLC using arbitrary non-punctual intervals) which are decidable in EXPSpace are expressible in TPTL^{0,∞}.

Finally, we leave open (i) the extension of this work with Past modalities, (ii) FOL-like characterizations of TPTL^{0,∞}, and (iii) whether adding multiple clocks in TPTL^{0,∞} improves expressiveness.

References

- 1 R. Alur, T. Feder, and T. Henzinger. The benefits of relaxing punctuality. *J.ACM*, 43(1):116–146, 1996.
- 2 Rajeev Alur and David L. Dill. A theory of timed automata. *Theor. Comput. Sci.*, 126(2):183–235, 1994. doi:10.1016/0304-3975(94)90010-8.

- 3 Rajeev Alur, Tomás Feder, and Thomas A. Henzinger. The benefits of relaxing punctuality. In Luigi Logrippo, editor, *Proceedings of the Tenth Annual ACM Symposium on Principles of Distributed Computing, Montreal, Quebec, Canada, August 19-21, 1991*, pages 139–152. ACM, 1991. doi:10.1145/112600.112613.
- 4 Rajeev Alur and Thomas A. Henzinger. Back to the future: Towards a theory of timed regular languages. In *33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, 24-27 October 1992*, pages 177–186. IEEE Computer Society, 1992. doi:10.1109/SFCS.1992.267774.
- 5 Rajeev Alur and Thomas A. Henzinger. Real-time logics: Complexity and expressiveness. *Inf. Comput.*, 104(1):35–77, 1993. doi:10.1006/inco.1993.1025.
- 6 Rajeev Alur and Thomas A. Henzinger. A really temporal logic. *J. ACM*, 41(1):181–203, January 1994. doi:10.1145/174644.174651.
- 7 Patricia Bouyer, Fabrice Chevalier, and Nicolas Markey. On the expressiveness of tptl and mtl. In Sundar Sarukkai and Sandeep Sen, editors, *FSTTCS 2005: Foundations of Software Technology and Theoretical Computer Science*, pages 432–443, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- 8 Thomas Ferrère. The compound interest in relaxing punctuality. In Klaus Havelund, Jan Peleska, Bill Roscoe, and Erik P. de Vink, editors, *Formal Methods – 22nd International Symposium, FM 2018, Held as Part of the Federated Logic Conference, FloC 2018, Oxford, UK, July 15-17, 2018, Proceedings*, volume 10951 of *Lecture Notes in Computer Science*, pages 147–164. Springer, 2018. doi:10.1007/978-3-319-95582-7_9.
- 9 Paul Gastin and Denis Oddoux. LTL with past and two-way very-weak alternating automata. In Branislav Rován and Peter Vojtás, editors, *Mathematical Foundations of Computer Science 2003, 28th International Symposium, MFCS 2003, Bratislava, Slovakia, August 25-29, 2003, Proceedings*, volume 2747 of *Lecture Notes in Computer Science*, pages 439–448. Springer, 2003. doi:10.1007/978-3-540-45138-9_38.
- 10 Christoph Haase, Joël Ouaknine, and James Worrell. On process-algebraic extensions of metric temporal logic. In A. W. Roscoe, Clifford B. Jones, and Kenneth R. Wood, editors, *Reflections on the Work of C. A. R. Hoare*, pages 283–300. Springer, 2010. doi:10.1007/978-1-84882-912-1_13.
- 11 Moritz Hammer, Alexander Knapp, and Stephan Merz. Truly on-the-fly ltl model checking. In Nicolas Halbwachs and Lenore D. Zuck, editors, *Tools and Algorithms for the Construction and Analysis of Systems*, pages 191–205, Berlin, Heidelberg, 2005. Springer Berlin Heidelberg.
- 12 Thomas A. Henzinger. It’s about time: Real-time logics reviewed. In Davide Sangiorgi and Robert de Simone, editors, *CONCUR’98 Concurrency Theory*, pages 439–454, Berlin, Heidelberg, 1998. Springer Berlin Heidelberg.
- 13 Thomas A. Henzinger, Jean-François Raskin, and Pierre-Yves Schobbens. The regular real-time languages. In Kim Guldstrand Larsen, Sven Skyum, and Glynn Winskel, editors, *Automata, Languages and Programming, 25th International Colloquium, ICALP’98, Aalborg, Denmark, July 13-17, 1998, Proceedings*, volume 1443 of *Lecture Notes in Computer Science*, pages 580–591. Springer, 1998. doi:10.1007/BFb0055086.
- 14 Y. Hirshfeld and A. Rabinovich. An expressive temporal logic for real time. In *MFCS*, pages 492–504, 2006.
- 15 Yoram Hirshfeld and Alexander Rabinovich. Expressiveness of metric modalities for continuous time. In Dima Grigoriev, John Harrison, and Edward A. Hirsch, editors, *Computer Science – Theory and Applications*, pages 211–220, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- 16 Hsi-Ming Ho. Revisiting timed logics with automata modalities. In Necmiye Ozay and Pavithra Prabhakar, editors, *Proceedings of the 22nd ACM International Conference on Hybrid Systems: Computation and Control, HSCC 2019, Montreal, QC, Canada, April 16-18, 2019*, pages 67–76. ACM, 2019. doi:10.1145/3302504.3311818.

- 17 James Jerson Ortiz, Axel Legay, and Pierre-Yves Schobbens. Memory event clocks. In Krishnendu Chatterjee and Thomas A. Henzinger, editors, *Formal Modeling and Analysis of Timed Systems*, pages 198–212, Berlin, Heidelberg, 2010. Springer Berlin Heidelberg.
- 18 S. N. Krishna K. Madnani and P. K. Pandya. On unary fragments of mtl over timed words. In *ICTAC*, pages 333–350, 2014.
- 19 Shankara Narayanan Krishna, Khushraj Madnani, Manuel Mazo Jr., and Paritosh K. Pandya. Generalizing non-punctuality for timed temporal logic with freeze quantifiers. In Marieke Huisman, Corina S. Pasareanu, and Naijun Zhan, editors, *Formal Methods – 24th International Symposium, FM 2021, Virtual Event, November 20-26, 2021, Proceedings*, volume 13047 of *Lecture Notes in Computer Science*, pages 182–199. Springer, 2021. doi:10.1007/978-3-030-90870-6_10.
- 20 Shankara Narayanan Krishna, Khushraj Madnani, and Paritosh K. Pandya. Logics meet 1-clock alternating timed automata. In Sven Schewe and Lijun Zhang, editors, *29th International Conference on Concurrency Theory, CONCUR 2018, September 4-7, 2018, Beijing, China*, volume 118 of *LIPICs*, pages 39:1–39:17. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2018. doi:10.4230/LIPICs.CONCUR.2018.39.
- 21 Slawomir Lasota and Igor Walukiewicz. Alternating timed automata. *ACM Trans. Comput. Log.*, 9(2):10:1–10:27, 2008. doi:10.1145/1342991.1342994.
- 22 Christof Loding and Wolfgang Thomas. Alternating automata and logics over infinite words. In Jan van Leeuwen, Osamu Watanabe, Masami Hagiya, Peter D. Mosses, and Takayasu Ito, editors, *Theoretical Computer Science: Exploring New Frontiers of Theoretical Informatics*, pages 521–535, Berlin, Heidelberg, 2000. Springer Berlin Heidelberg.
- 23 Khushraj Nanik Madnani. *On Decidable Extensions of Metric Temporal Logic*. PhD thesis, Indian Institute of Technology Bombay, Mumbai, India, 2019.
- 24 J. Ouaknine and J. Worrell. On the decidability of metric temporal logic. In *LICS*, pages 188–197, 2005.
- 25 J. Ouaknine and J. Worrell. Safety metric temporal logic is fully decidable. In *TACAS*, pages 411–425, 2006.
- 26 Paritosh K. Pandya and Simoni S. Shah. On expressive powers of timed logics: Comparing boundedness, non-punctuality, and deterministic freezing. In Joost-Pieter Katoen and Barbara König, editors, *CONCUR 2011 – Concurrency Theory – 22nd International Conference, CONCUR 2011, Aachen, Germany, September 6-9, 2011. Proceedings*, volume 6901 of *Lecture Notes in Computer Science*, pages 60–75. Springer, 2011. doi:10.1007/978-3-642-23217-6_5.
- 27 Paritosh K. Pandya and Simoni S. Shah. The unary fragments of metric interval temporal logic: Bounded versus lower bound constraints. In *Automated Technology for Verification and Analysis – 10th International Symposium, ATVA 2012, Thiruvananthapuram, India, October 3-6, 2012. Proceedings*, pages 77–91, 2012.
- 28 Radek Pelánek and Jan Strejček. Deeper connections between ltl and alternating automata. In Jacques Farré, Igor Litovsky, and Sylvain Schmitz, editors, *Implementation and Application of Automata*, pages 238–249, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- 29 A. Rabinovich. Complexity of metric temporal logic with counting and pnueli modalities. In *FORMATS*, pages 93–108, 2008.
- 30 Alexander Rabinovich. Complexity of metric temporal logics with counting and the pnueli modalities. *Theor. Comput. Sci.*, 411(22-24):2331–2342, 2010. doi:10.1016/j.tcs.2010.03.017.
- 31 Jean Francois Raskin. *Logics, Automata and Classical Theories for Deciding Real Time*. PhD thesis, Universite de Namur, 1999.
- 32 Gareth Scott Rohde. *Alternating Automata and the Temporal Logic of Ordinals*. PhD thesis, University of Illinois at Urbana-Champaign, USA, 1997. AAI9812757.
- 33 Heikki Tauriainen. *Automata and linear temporal logic: Translations with transition-based acceptance*. Doctoral thesis, Helsinki University of Technology, 2006.

23:18 Satisfiability Checking for $TPTL^{0,\infty}$ Is PSPACE-Complete

- 34 Moshe Y. Vardi. An automata-theoretic approach to linear temporal logic. In Faron Moller and Graham Birtwistle, editors, *Logics for Concurrency: Structure versus Automata*, pages 238–266, Berlin, Heidelberg, 1996. Springer Berlin Heidelberg. doi:10.1007/3-540-60915-6_6.
- 35 Thomas Wilke. Specifying timed state sequences in powerful decidable logics and timed automata. In *Formal Techniques in Real-Time and Fault-Tolerant Systems, Third International Symposium Organized Jointly with the Working Group Provably Correct Systems – ProCoS, Lübeck, Germany, September 19-23, Proceedings*, pages 694–715, 1994. doi:10.1007/3-540-58468-4_191.