

May 1991

Report No. STAN-CS-91-1359

The Benefits of Relaxing Punctuality


by

R. Alur, T. Feder, and T. Henzinger

Department of Computer Science

Stanford University

Stanford, California 94305



REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE 5/28/91	3. REPORT TYPE AND DATES COVERED
----------------------------------	---------------------------	----------------------------------

4. TITLE AND SUBTITLE THE BENEFITS OF RELAXING PUNCTUALITY	5. FUNDING NUMBERS
---	--------------------

6. AUTHOR(S) RAJEEV ALUR, TOMÁS FEDER, THOMAS A. HENZINGER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) DEPT. OF COMPUTER SCIENCE STANFORD UNIVERSITY STANFORD, CA 94305	8. PERFORMING ORGANIZATION REPORT NUMBER
--	--

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) DARPA ARLINGTON, VA 22209	10. SPONSORING / MONITORING AGENCY REPORT NUMBER N00039-84C-0211
---	---

II. SUPPLEMENTARY NOTES

12a. DISTRIBUTION / AVAILABILITY STATEMENT unlimited	12 b. DISTRIBUTION CODE
---	-------------------------

13. ABSTRACT (Maximum 200 words)

Abstract. The most natural, compositional way of modeling real-time systems uses a dense domain for time. The satisfiability of real-time constraints that are capable of expressing punctuality in this model is, however, known to be undecidable.

We introduce a temporal language that can constrain the time difference between events only with finite (yet arbitrary) precision and show the resulting logic to be EXPSPACE-complete. This result allows us to develop an algorithm for the verification of timing properties of real-time systems with a dense semantics.

14. SUBJECT TERMS	15. NUMBER OF PAGES 38
	16. PRICE CODE

17. SECURITY CLASSIFICATION OF REPORT	18. SECURITY CLASSIFICATION OF THIS PAGE	19. SECURITY CLASSIFICATION OF ABSTRACT	20. LIMITATION OF ABSTRACT
---------------------------------------	--	---	----------------------------

The Benefits of Relaxing Punctuality^{*†}

Rajeev Alur[‡] Tomás Feder[§] Thomas A. Henzinger[†]

May 29, 1991

Abstract. The most natural, compositional way of modeling real-time systems uses a dense domain for time. The satisfiability of real-time constraints that are capable of expressing punctuality in this model is, however, known to be undecidable.

We introduce a temporal language that can constrain the time difference between events only with finite (yet arbitrary) precision and show the resulting logic to be EXPSPACE-complete. This result allows us to develop an algorithm for the verification of timing properties of real-time systems with a dense semantics.

1 Introduction

The formal study of reactive systems has led recently to a number of suggestions of how real-time requirements of such systems ought to be modeled, specified, and verified. Most of these approaches are situated at either extreme of the trade-off between *realistic modeling* of time and *feasible verification* of timing properties. Typically, they either use a continuous model of time at the expense of decidability [ACD90, Koy90, Lew90], or they sacrifice continuity to obtain decision procedures [JM86, AH89, AH90, EMSS89, HLP90, Ost90]. This paper shows how a slight relaxation of the notion of punctuality allows us to combine the best of both worlds.

*An abbreviated version of this paper appears in the proceedings of the *Tenth Annual ACM Symposium on Principles of Distributed Computing* (1991).

†This research was supported in part by an IBM graduate fellowship, by the National Science Foundation grants CCR-89-11512, CCR-89-13641, and MIP-H-588807, by the Defense Advanced Research Projects Agency under contract N00039-84-C-0211, and by the United States Air Force Office of Scientific Research under contract AFOSR-90-0057.

‡Department of Computer Science, Stanford University, Stanford, CA 94305.

§Bell Communications Research, Morristown, NJ 07962.

Let us be more specific. The linear (trace) semantics of a reactive system is defined as a set of possible behaviors, each of which is represented by a sequence of system states. This model is most naturally extended to incorporate real time by associating, with every state, an interval of the real line, which indicates the period of time during which the system is in that state. That is, we represent the possible behaviors of a real-time system by *timed state sequences*.

Alas, even the satisfiability of a very simple class of real-time properties turns out to be undecidable in this model [AH89]. An inspection of the proof shows that the only timing constraints required are of the form

$$\Box \Box \rightarrow \Diamond_{=5} q), \quad (\dagger)$$

predicting that every p-state is followed by a q-state precisely 5 time units later.

This negative result has led us, at first, to weaken the expressiveness of the model by adopting the *semantic* abstraction that, at every state change, we may record only a discrete approximation — the number of ticks of a digital clock — to the real time. Thus we have interpreted the formula (\dagger) to require only that the p-state and the corresponding q-state are separated by exactly 5 clock ticks; their actual difference in time may be as much as (say) 5.9 time units or as small as 4.1 time units. We have shown that several interesting real-time logics are decidable under this weaker, *digital-clock*, interpretation [AH89, AH90].

In this paper we pursue an alternative, *syntactic*, concession. Instead of digitizing the meaning of a sentence, we prohibit timing constraints that predict the time difference between two states with infinite accuracy. In particular, we may not state the property given above, but only an approximation such as

$$\Box (p \rightarrow \Diamond_{(4.9,5.1)} q),$$

requiring that the p-state and the corresponding q-state are separated by more than 4.9 time units and less than 5.1 time units.

We define a language that can constrain the time difference between events only with finite (yet arbitrary) precision. The resulting *metric interval temporal logic* MITL is shown to be decidable in EXPSPACE. Furthermore, we show how to verify a real-time system with respect to a specification in MITL.

Properties of timed state sequences can, alternatively, be defined by *timed automata* [AD90]. While the emptiness problem for these automata

is solvable, they are not closed under complement. MITL identifies a fragment of the properties definable by timed automata that is closed under all boolean operations. Thus the novelty of our results is that they give a logical formalism with a *continuous* interpretation of time that is suitable for the automatic verification and synthesis of finite-state real-time systems.

Both the semantic abstraction of digitizing models as well as the syntactic restriction of excluding equality in timing constraints limit the real-time properties that are definable in a similar way: they rule out the notion of absolute punctuality and replace it by a looser concept of almost-on-time behavior. This sacrifice is viable because, by choosing the clock tick of the digital clock small enough, we can still achieve arbitrary precision in either approach; moreover, the corresponding costs for achieving the desired accuracy are the same.

Yet the introduction of a mandatory slack through the syntax (rather than through the semantics) turns out to be the more powerful technique: we show that the properties of timed state sequences that can be defined in MITL are a proper superset of those definable with equality under a digital-clock interpretation. Also, many of the practically interesting forms of punctuality are still expressible in MITL, such as the requirement that every p-state is separated from the *closest* subsequent q-state by precisely 5 time units.

The remainder of the paper is organized in four parts. In Section 2, we introduce and motivate the logic MITL, and show it to be more expressive than digitization. In Section 3, we introduce a variant of timed automata as a model for finite-state real-time systems. In Section 4, we reduce the decision problem for MITL to the emptiness problem of timed automata. In the concluding section, we show how the results of this paper lead to an algorithm that verifies MITL-specifications of real-time systems that are given as timed automata.

We remark that in this paper we introduce MITL with future temporal operators only. All of our results, in particular EXPSPACE-completeness, generalize to MITL with both future and *past* temporal operators.

2 Metric Interval Temporal Logic

We define timed state sequences as formal models of real-time behavior. Then we introduce a temporal language to define properties of timed state sequences and study its expressive power.

2.1 Intervals and interval sequences

An interval is a convex subset of the nonnegative real numbers \mathbb{R}^+ . Intervals may be open, halfopen, or closed; bounded or unbounded. More precisely, each interval is of one of the following forms: $[a, b]$, $[a, b)$, $[a, \infty)$, $(a, b]$, (a, b) , (a, ∞) , where $a \leq b$ and $a, b \in \mathbb{R}^+$. For an interval I of the above form, a is its left end-point, and b is its right end-point; the left end-point of I is denoted by $l(I)$ and the right end-point, for bounded I , is denoted by $r(I)$.

An interval I is *singular* iff it is of the form $[a, a]$; that is, I is closed and $l(I) = r(I)$.

Two intervals I and I' are *adjacent* iff (1) either I is right-open and I' is left-closed, or I is right-closed and I' is left-open, and (2) $r(I) = l(I')$. For instance, the intervals $(1, 2]$ and $(2, 2.5)$ are adjacent.

An *interval sequence* $\tau = I_0 I_1 I_2 I_3 \dots$ is a finite or infinite sequence of intervals that partitions \mathbb{R}^+ :

1. Any two neighboring intervals I_i and I_{i+1} are adjacent.
2. For all $t \in \mathbb{R}^+$, there is some interval I_i with $t \in I_i$.

In particular, I_0 is left-closed and $l(I_0) = 0$; if τ is finite, then its last interval must be unbounded.

We will freely use intuitive pseudo-arithmetic expressions to denote intervals. For example, the expressions $\leq b$ and $> a$ stand for the intervals $[0, b]$ and (a, ∞) , respectively; by $< I$ we denote the interval $\{t' \mid 0 \leq t' < t \text{ for all } t \in I\}$. The expression $t + I$, where I is an interval and $t \in \mathbb{R}^+$, denotes the interval $\{t + t' \mid t' \in I\}$; similarly, the expressions $I - t$ and $t I$ stand for the intervals $\{t' - t \mid t' \in I \text{ and } t' \geq t\}$ and $\{tt' \mid t' \in I\}$, respectively.

2.2 Timed state sequences

Let P be a finite set of atomic propositions. We assume that, at any point in time, the global state of a (finite-state) system can be modeled by an interpretation (or truth-value assignment) for P . We therefore identify states s with subsets of P ; that is, $s \models p$ iff $p \in s$ (for $p \in P$).

A behavior of a discrete system over time can, consequently, be modeled by a finite or infinite sequence

$$\rho: (s_0, I_0) \rightarrow (s_1, I_1) \rightarrow (s_2, I_2) \rightarrow (s_3, I_3) \rightarrow \dots$$

of states $s_i \in 2^P$ and corresponding time intervals $I_i \subseteq \mathbb{R}^+$. A *timed state sequence* $\rho = (\sigma, \tau)$ consists of a sequence $\sigma : s_0 s_1 s_2 \dots$ of states and an interval sequence $\tau : I_0 I_1 I_2 \dots$ of the same length.

A timed state sequence $\rho = (\sigma, \tau)$ can be viewed as a map ρ^* from the time domain \mathbb{R}^+ to the states 2^P (let $\rho^*(t) = s_i$ if $t \in I_i$). Thus a timed state sequence provides complete information about the global state of a system at each time instant: at time $t \in I_i$, the system is in state $\rho^*(t) = s_i$. Timed state sequences obey the *finite-variability* condition: between any two points in time there are only finitely many state changes. This assumption is adequate for modeling *discrete* systems.

Given a timed state sequence (σ, τ) , the i -th transition point, denoted by t_i , is defined to be the left end-point of the interval I_i ; that is, $t_i = l(I_i)$. Note that the state at time t_i is s_{i-1} if I_i is left-open, and is s_i if I_i is left-closed.

Our definition allows *transient* states, which occur only at a single point in time. If I_i is a singular interval $[t_i, t_i]$, then the state at time t_i is s_i , but the state just before t_i is s_{i-1} , and the state just after t_i is s_{i+1} . Observe that in such a case neither s_{i-1} nor s_{i+1} can be transient, because the interval I_{i-1} must be right-open and the interval I_{i+1} must be left-open. Transient states are useful for modeling the truth of propositions that represent instantaneous events and, thus, are true only at isolated points in time.

We will also need the concept of a *suffix* of a timed state sequence. For a timed state sequence $\rho = (\sigma, \tau)$ and time $t \in I_i$, let $\rho^t = (\sigma^t, \tau^t)$ be the timed state sequence with the state component $\sigma^t : s_i s_{i+1} s_{i+2} \dots$ and the time component

$$\tau^t : (I_i - t)(I_{i+1} - t)(I_{i+2} - t) \dots$$

Note that the suffix operator is defined such that $(\rho^t)^*(t') = \rho^*(t + t')$ for all $t' \in \mathbb{R}^+$. In particular, $\rho^0 = \rho$.

2.3 Syntax and semantics of MITL

We introduce an extension of linear temporal logic, *metric interval temporal logic* (or MITL), that is interpreted over *timed* state sequences. A standard way of adding timing requirements to temporal languages is to replace the temporal operators with time-constrained versions, such as the constrained *eventually* operator $\diamond_{[2,4]}$ meaning “eventually within 2 to 4 time units” [EMSS89, AH90, Koy90]. We adopt this approach for MITL, with the restriction that operators cannot be constrained by singular time intervals.

The formulas of MITL are built from atomic propositions by boolean connectives and time-constrained versions of the *until* operator \mathcal{U} ; they are defined inductively as follows:

$$\phi := p \mid \neg\phi \mid \phi_1 \wedge \phi_2 \mid \phi_1 \mathcal{U}_I \phi_2,$$

where $p \in P$ and I is a *nonsingular* interval with rational end-points (I may be unbounded).

The formulas of MITL are interpreted over timed state sequences, which provide an interpretation for the atomic propositions at each time instant. Informally, the formula $\phi_1 \mathcal{U}_I \phi_2$ holds at time $t \in \mathbb{R}^+$ of a timed state sequence iff there is a later time instant $t' \in t + I$ such that ϕ_2 holds at time t' and ϕ_1 holds throughout the time interval (t, t') .

Given an MITL-formula ϕ and a timed state sequence $\rho = (a, \tau)$, the satisfaction relation $\rho \models \phi$ is defined inductively as follows:

$$\begin{aligned} \rho \models p &\text{ iff } p \in s_0. \\ \rho \models \neg\phi &\text{ iff } \rho \not\models \phi. \\ \rho \models \phi_1 \wedge \phi_2 &\text{ iff } \rho \models \phi_1 \text{ and } \rho \models \phi_2. \\ \rho \models \phi_1 \mathcal{U}_I \phi_2 &\text{ iff } \rho^t \models \phi_2 \text{ for some } t \in I, \text{ and } \rho^{t'} \models \phi_1 \text{ for all } \\ &t' \in (0, t). \end{aligned}$$

The MITL-formula ϕ is *satisfiable (valid)* iff $\rho \models \phi$ for some timed state sequence ρ (all timed state sequences ρ , respectively).

Observe that the logic MITL is insensitive to *stuttering*. Given two timed state sequence $\rho = (\sigma, \tau)$ and $\rho' = (\sigma', \tau')$ such that ρ' has a subsequence of the form

$$(s_{i-1}, I_{i-1}) \rightarrow (s_i, I) \rightarrow (s_i, I') \rightarrow (s_{i+1}, I_{i+1})$$

and $I \cup I' = I_i$, then $\rho^* = \rho'^*$, and $\rho \models \phi$ iff $\rho' \models \phi$ for every MITL-formula ϕ .

The satisfaction relation has another desirable property: the truth value of any MITL-formula does not change more than ω times along a timed state sequence. Thus timed state sequences satisfy the finite-variability condition not only with respect to the truth of atomic propositions, but also with respect to arbitrarily complex MITL-formulas. The following lemma states this property formally:

Lemma 2.1 (Model refinement) *Let ϕ be an MITL-formula and $\rho = (\sigma, \tau)$ be a timed state sequence. There exists an interval sequence τ_ϕ :*

$J_0 J_1 \dots$ such that whenever t and t' belong to the same interval J_i , we have $\rho^t \models \psi$ iff $\rho^{t'} \models \psi$ for each subformula ψ of ϕ . Moreover, if all interval end-points in τ are rational numbers, then so are all interval end-points in τ_ϕ .

Proof of Lemma 2.1 Let $\rho = (a, \tau)$. The proof is by induction on the structure of ϕ . For an atomic proposition p , take τ_p to be τ . For a negated formula $\neg\phi'$, take $\tau_{\neg\phi'}$ to be $\tau_{\phi'}$. In case of the conjunction $\phi_1 \wedge \phi_2$, the interval sequence $\tau_{\phi_1 \wedge \phi_2}$ is constructed by taking the intersection of the two interval sequences τ_{ϕ_1} and τ_{ϕ_2} .

Now let us consider the case that ϕ has the form $\phi_1 \mathcal{U}_I \phi_2$. Let $\tau_{\phi_1 \wedge \phi_2}$ be the interval sequence $J_0 J_1 \dots$. We construct a refinement $\tau_\phi : J'_0 J'_1 \dots$ of $\tau_{\phi_1 \wedge \phi_2}$ such that whenever t and t' are in the same interval J'_i , then both t and t' belong to the same interval J_k , both $t + Z(1)$ and $t' + Z(1)$ belong to the same interval J_l , and, if I is bounded, both $t + r(1)$ and $t' + r(1)$ belong to the same interval J_m , for some k, l, m . It is clear that such a sequence can be constructed by a finite splitting of each interval J_i such that, if the end-points of all intervals J_i are rational, then so are the end-points of all intervals J'_i . Furthermore, it is **easy** to check that $\rho^t \models \phi$ iff $\rho^{t'} \models \phi$ whenever t and t' are in the same interval J'_i . ■

For any MITL-formula ϕ , we say that the timed state sequence $\rho = (a, \tau_\phi)$ is *ϕ -fine*. Clearly, ϕ is satisfiable iff it has a &fine model.

2.4 Defined operators

Now let us introduce some standard abbreviations for additional temporal operators. The defined operators $\diamond_I \phi$ (constrained *eventually*) and $\square_I \phi$ (constrained *always*) stand for $\text{true } \mathcal{U}_I \phi$ and $\neg \diamond_I \neg \phi$, respectively. It follows that the formula $\square_I \phi$ (or $\diamond_I \phi$) holds at time $t \in \mathbb{R}^+$ of a timed state sequence iff ϕ holds at all times (at some time, respectively) within the interval $t + I$.

We usually suppress the interval $(0, \infty)$ as a subscript. Thus the MITL-operators \diamond , \square , \mathcal{U} coincide with the conventional unconstrained *strict eventually*, *strict always*, and *strict until* operators of temporal logic. This is because the *until* operator of MITL is implicitly strict in its first argument. The corresponding non-strict operators are definable in MITL as $\diamond_{[0, \infty)}$ (also written $\diamond_{>0}$), $\square_{[0, \infty)}$, and $\mathcal{U}_{>0}$, and

$$\phi_2 \vee (\phi_1 \wedge \phi_1 \mathcal{U} \phi_2)$$

for $\phi_2 \mathcal{U}^= \phi_1$ (where $\mathcal{U}^=$ denotes the unconstrained non-strict *until* operator). Note that, on the other hand, the operator \mathcal{U}_I cannot be defined in terms of an *until* operator that is not strict in its first argument; this is why we have chosen the *strict* versions of temporal operators to be primitive.

Using these abbreviations, the typical bounded response property that “every p-state is followed by a q-state within 5 time units,” can be expressed by the MITL-formula

$$\Box_{\geq 0} (p \rightarrow \Diamond_{(0,5]} q).$$

We also define a constrained *unless* operator as the dual of the *until* operator:

$$\phi_1 \mathcal{I} \mathcal{U} \phi_2 \text{ stands for } \neg((\neg\phi_2) \mathcal{U}_I (\neg\phi_1)).$$

It follows that the formula $\phi_1 \mathcal{I} \mathcal{U} \phi_2$ holds at time $t \in \mathbb{R}^+$ of a timed state sequence iff either ϕ_1 is true throughout the interval $t + I$, or there is a time instant $t' > t$ such that ϕ_2 is true at time t' and ϕ_1 holds at all instants $t'' \leq t'$ within the interval $t + I$. Note that the unconstrained version $\phi_1 \mathcal{U} \phi_2$ of the *unless* operator of MITL differs slightly from the conventional strict *unless* operator, which can be defined as $\phi_1 \mathcal{U} (\phi_1 \wedge \phi_2)$.

We can apply the definition of the *unless* operator to move negations through *until* operators. Thus we may obtain, from any MITL-formula, an equivalent formula, containing both *until* and *unless* operators, in which all negations are in front of atomic propositions.

2.5 Avoiding undecidability

A few comments on our choice of syntax are in order. First, MITL has no *next-time* operator, because due to the density of the time domain there is no unique next time. Also, MITL is, syntactically viewed, essentially the restriction of *metric temporal logic* (MTL [AH90]) that prohibits the use of equality in time bounds. For example, in MITL we cannot directly express the punctuality condition that “every p-state is followed by a q-state after exactly 5 time units,”

$$\Box \otimes \Box (p \rightarrow \Diamond_{=5} q),$$

because the singular interval $[5, 5]$ is not allowed as a subscript. We will show that there is, in fact, no MITL-formula that expresses this condition, and that the restriction of MITL to *nonsingular* intervals is essential for decidability.

Note that some practically important forms of equality are expressible in MITL; we define $(\neg\phi) \mathcal{U}_{=n} \phi$, for $n > 0$, as an abbreviation for the MITL-formula $\Box_{(0,n)} \neg\phi \wedge \Diamond_{(0,n]} \phi$. Thus the stronger condition that “for every p-state the densest subsequent q-state is after exactly 5 time units,”

$$\Box_{\geq 0} (p \rightarrow (\neg q) \mathcal{U}_{=5} q),$$

is expressible in MITL.

Let $\text{MITL}_{=}$ be the extension of MITL that admits singular intervals as time bounds on the temporal operators. We show that the decision problem of $\text{MITL}_{=}$ is complete for the complexity class Π_1^1 , which is situated in the analytical hierarchy strictly above all recursively enumerable sets (see, for example, [Rog67]). It follows that $\text{MITL}_{=}$ is not even axiomatizable.

Theorem 2.1 (MITL with equality) *The decision problem of $\text{MITL}_{=}$ is Π_1^1 -complete.*

Proof of Theorem 2.1 [Hi-hardness] The decision problem for *dense* MTL is Hi-complete [AH90]. A close inspection of the proof given there reveals that that only one operator with a singular subscript, $\Diamond_{=n}$ for any $n > 0$, is used to demonstrate Hi-hardness.

There is, however, a subtle difference between the dense interpretations defined in [AH90] and timed state sequences: a dense interpretation consists of an infinite sequence of states and corresponding time *instants*, not intervals. Consequently, while the formula $\Box_I \text{false}$ (for any finite nonempty interval I) is not satisfiable by any timed sequence, it is satisfiable by infinitely many dense interpretations — those that do not contain any states with times in I .

With some care we can still reduce the decision problem for dense MTL to the decision problem for MITL with equality, which demonstrates the Π_1^1 -hardness of the latter logic. Let r be a proposition that is true in infinitely many transient states and nowhere else; that is,

$$\phi_r: \Box \# \Box \mathcal{U}_{=1} \# \Box \rightarrow (\neg r) \mathcal{U} r).$$

It is not hard to see that a dense MTL-formula ϕ is valid iff the MITL-formula $\phi_r \rightarrow \phi^*$ is valid, where ϕ^* is obtained from ϕ by replacing every occurrence of a subformula $\psi_1 \mathcal{U}_I \psi_2$ with

$$(r \rightarrow \psi_1) \mathcal{U}_I (r \wedge \psi_2).$$

[Containment in Π_1^1] We show that the validity of a formula ϕ of MITL₌ can be phrased as a Π_1^1 -sentence, asserting that all timed state sequences are models of ϕ . From Theorem 2.2 to be proved shortly, it follows that if ϕ has a model, then it has a model in which all interval end-points are rational numbers (i.e., a *rational* model). This observation allows us to assert the validity of ϕ as a Π_1^1 -sentence: ϕ is valid iff $\rho \models \phi$ for all rational models ρ . It is routine to encode a rational model by a set of natural numbers, and to express the satisfaction relation in first-order arithmetic. ■

Another possible extension of the syntax of MITL is to permit time bounds on *both* arguments of the *until* operator, as is the case for all logics that admit explicit references to time in atomic formulas (such as TPTL [AH89]). The intended meaning of the formula $\phi_1 \mathcal{U}_I \phi_2$ at time $t \in \mathbb{R}^+$ of a timed state sequence is that there is a later time instant $t' \in t + I$ such that ϕ_2 holds at time t' and ϕ_1 holds throughout the time interval $(t + I) \cap [t, t']$. Such an extension leads, however, again to undecidability. This is because the role of $\diamond_{=n} \phi$ in the undecidability argument for MITL, can be replaced by the formula $false \geq_n \mathcal{U}_{\geq n} \phi$.

2.6 Real versus rational time

Having justified our choice of syntax, let us look at other options for defining the semantics of MITL. While timed state sequences are defined by choosing the set of (nonnegative) reals to model time, for interpreting formulas of MITL, the crucial property of the time domain \mathbb{R}^+ is not its continuity, but only its denseness. In particular, we show that replacing the time domain \mathbb{R}^+ with the nonnegative rational numbers \mathbb{Q}^+ when defining the semantics of MITL does not change the satisfiability (and validity) of any MITL-formula.

We call a timed state sequence (σ, τ) *rational* iff the end-points of all intervals in τ are rational. A formula ϕ of MITL₌ is said to be *Q-satisfiable* iff $\rho \models \phi$ for some rational timed state sequence ρ , where the satisfaction relation \models is redefined so that all time quantifiers range over \mathbb{Q}^+ only.

We show that this new notion of satisfiability is the same as the old one. In other words, MITL-formulas cannot distinguish the time domain \mathbb{R}^+ from the time domain \mathbb{Q}^+ . This equivalence of real and rational models follows from the following two lemmas.

Lemma 2.2 (Rational models) *Let ϕ be an MITL-formula and ρ a rational ϕ -fine timed state sequence. Then ρ Q-satisfies ϕ iff $\rho \models \phi$.*

Proof of Lemma 2.2 We use induction on the structure of ϕ . Let us consider only the interesting case, that ϕ has the form $\phi_1 \mathcal{U}_I \phi_2$.

Suppose that $\rho = (a, \tau)$ is rational, #-fine, and Q-satisfies ϕ ; that is, ρ^t Q-satisfies ϕ_2 for some rational $t \in I$, and $\rho^{t'}$ Q-satisfies ϕ_1 for all rationals $0 < t' < t$. By the induction hypothesis, we may conclude that $\rho^t \models \phi_2$ and $\rho^{t'} \models \phi_1$ for all rationals $0 < t' < t$. Since $t \in \mathbb{R}^+$, it remains to be shown that $\rho^{t''} \models \phi_1$ for all reals $0 < t'' < t$. Consider an arbitrary real $0 < t'' < t$, and assume that $t'' \in I$. Since ρ is rational, there is also a rational $t' \in I$; with $0 < t' < t$. We know that $\rho^{t'} \models \phi_1$ and, since ρ is \$-fine, it follows that $\rho^{t''} \models \phi_1$.

The second direction, that every rational ϕ -fine model of ϕ Q-satisfies ϕ , follows by a similar argument. ■

For any MITL-formula ϕ , let n_ϕ be the least common denominator of all (rational) interval end-points in ϕ ; that is, all constants in ϕ are multiples of $1/n_\phi$.

Lemma 2.3 (Model equivalence) *Let $\rho = (\sigma, \tau)$ and $\rho' = (a, \tau')$ be two timed state sequences, and ϕ be a formula of MITL₌. Suppose that for all $t \in \mathbb{R}^+$, if $t = t_i + m/n_\phi$ for some left end-point t_i of an interval in τ and some nonnegative integer $m \in \mathbb{N}$, then $t \in I_j$ iff $t \in I'_j$. Then $\rho \models \phi$ iff $\rho' \models \phi$.*

Proof of Lemma 2.3 We write $\rho \sim \rho'$ iff the two timed state sequences ρ and ρ' satisfy the premise of the lemma. First observe that, if $\rho \sim \rho'$ for $\rho = (\sigma, \tau)$ and $\rho' = (\sigma, \tau')$ and $t \in I$, then we can find $f(t) \in I'_i$ such that $\rho^t \sim \rho'^{f(t)}$. Furthermore, $f(t) < f(t')$ iff $t < t'$.

Using this observation, the lemma follows by straightforward induction on the structure of ϕ . ■

Lemma 2.3 classifies timed state sequences into equivalence classes such that the members of a class cannot be distinguished by formulas of MITL₌. It implies, in particular, the following theorem:

Theorem 2.2 (Rational time) *A formula ϕ of MITL₌ is Q-satisfiable iff it is satisfiable.*

Proof of Theorem 2.2 Suppose that ϕ is Q-satisfiable in the rational model p . By Lemma 2.1, there is a rational +-fine refinement of p that Q-satisfies ϕ . By Lemma 2.2, this refinement is a (real) model of ϕ .

The proof of the second direction uses Lemma 2.3. Consider a (real) model ρ of ϕ . The lemma allows us to adjust the interval boundaries in ρ

as long as (1) no interval is adjusted across multiples of $1/n_\phi$, and (2) the ordering of the fractional parts (modulo $1/n_\phi$) of all interval boundaries is not altered. The denseness of \mathbb{Q}^+ allows us to adjust all boundaries to be rational numbers. The resulting rational timed state sequence is a (real) model of ϕ and, by Lemma 2.1 and Lemma 2.2, its δ -refinement \mathbb{Q} -satisfies ϕ . ■

2.7 Expressive power of MITL

We define the semantics of a system as a set of timed state sequences; such a set is called a *real-time property*. Every formula ϕ of a real-time logic (say, MITL) specifies a real-time property — the set of models of ϕ . The expressive power of a logic is measured by the real-time properties that can be specified by formulas of the logic.

We compare the expressive power of MITL to the use of a digital clock and MTL, which admits singular intervals as time bounds on temporal operators. More precisely, we show that the analog-clock model without equality (MITL) is more expressive than any digital-clock model with equality (MTL).

First let us review the definition of the logic MTL [AH90]. The syntax of MTL is the same as that of MITL₌. The formulas of MTL are interpreted over observation sequences. An *observation sequence* ρ is an infinite sequence

$$(s_0, T_0) \rightarrow (s_1, T_1) \rightarrow (s_2, T_2) \rightarrow (s_3, T_3) \rightarrow \dots$$

of observations. Each *observation* consists of a state $s_i \in 2^P$ and a time stamp $T_i \in \mathbb{N}$. The observation sequence ρ satisfies the *initiality* condition that $T_0 = 0$, the *monotonicity* condition that $T_i \leq T_{i+1}$ for all $i \geq 0$, and the *progress* condition that, for all $n \in \mathbb{N}$, there is some $i \geq 0$ such that $T_i > n$.

For an observation sequence ρ and an MTL-formula ϕ , the satisfaction relation $\rho \models \phi$ is defined as usual by induction on the structure of ϕ . The following clause considers the case of the (strict) *until* operator:

$$\rho \models \phi_1 \mathcal{U}_I \phi_2 \text{ iff } \rho^i \models \phi_2 \text{ for some } i \geq 0 \text{ with } T_i \in I, \text{ and } \rho^j \models \phi_1 \\ \text{for all } 0 < j < i.$$

(For an observation sequence ρ and $i \in \mathbb{N}$, the observation sequence ρ^i is the suffix of the shifted sequence $\rho - T_i$ that begins with the observation $(s_i, 0)$.) We consider only the fragment of MTL without the *next-state* operator; this restriction makes MTL-formulas insensitive to stuttering.

We need to formalize which real-time properties can be specified in MTL. To this end, let us consider how to extract an observation sequence from a timed state sequence ρ that describes the actual behavior of a real-time system. Observations are made with respect to a digital clock; the observation at time t records the state $\rho(t)$ and the value of the clock at time t . Clearly the observations depend on how fast the clock ticks, and at what time the clock is started.

Consequently, we define a *digital clock* $D = (\delta, \epsilon)$ to be a pair consisting of the distance $\delta \in \mathbb{R}^+$ between two successive clock ticks and the time $\epsilon \in \mathbb{R}^+$ of the first clock tick; that is, $0 \leq \epsilon < \delta$. At time $t \in \mathbb{R}^+$ the clock D shows the integer value $t_D = \lceil (t - \epsilon) / \delta \rceil$. The clock D is called *rational* iff both δ and ϵ are rational numbers.

The D -observation of the timed state sequence ρ at time t is $O_t = (\rho^*(t), t_D)$. As time increases, the D -observation stays the same until either the clock ticks or the state changes along ρ . All possible D -observations along ρ can be described by an w -sequence: the D -observed behavior of ρ is the observation sequence

$$\rho_D: O_{t_0} \rightarrow O_{t_1} \rightarrow O_{t_2} \rightarrow \dots,$$

such that for all $i \geq 0$, (1) $t_i < t_{i+1}$, and (2) for all $t \in (t_i, t_{i+1})$, O_t equals either O_{t_i} or $O_{t_{i+1}}$. These properties define ρ_D uniquely modulo stuttering (i.e., duplication of neighboring observations). Furthermore, the state component of ρ_D is the state component of ρ (modulo stuttering) with, if ρ is finite, infinite repetition of the final state.

For instance, consider the timed state sequence ρ :

$$(s_0, [0, 1)) \rightarrow (s_1, [1, 1]) \rightarrow (s_2, (1, 1.5]) \rightarrow (s_3, (1.5, \infty)).$$

Then the digital clock $(1, 0.5)$ observes the observation sequence $\rho_{(1, 0.5)}$:

$$\begin{aligned} (s_0, 0) &\rightarrow (s_0, 1) \rightarrow (s_1, 1) \rightarrow (s_2, 1) \rightarrow \\ (s_3, 2) &\rightarrow (s_3, 3) \rightarrow (s_3, 4) \rightarrow \dots \end{aligned}$$

For every digital clock D , every formula ϕ of MTL specifies a real-time property Π_ϕ^D — the set of timed state sequences ρ such that $\rho_D \models \phi$. We say that the MTL-formula ϕ *D-specifies* the real-time property Π_ϕ^D .

Now we can be specific about the sense in which the analog-clock model is, even without equality, more expressive than the digital-clock model, for any choice of digital clock.

Theorem 2.3 (Expressiveness of MITL) (a) Every real-time property that can be D -specified by an MTL-formula for some rational digital clock D , can also be specified in MITL. (b) There is a real-time property that can be specified in MITL but not D -specified by any MTL-formula for any digital clock D .

Proof of Theorem 2.3 (a) Given a rational clock $D = (\delta, \epsilon)$ and a formula ϕ of MTL, we construct an MITL-formula that specifies the real-time property $\Pi_{\phi}^{\mathcal{L}^D}$. We assume that ϕ contains only intervals of the form $[0, 0]$, $[1, 1]$, $[m, n]$ for $2 \leq m \leq n$, and $[m, \infty)$ for $m \geq 2$. It is trivial to convert any MTL-formula into this form; for instance, the MTL-formula $\diamond_{<5} \psi$ is equivalent to the formula $\diamond_{=0} \psi \vee \diamond_{=1} \psi \vee \diamond_{[2,4]} \psi$.

We model the ticks of the digital clock D by a new proposition r that holds only in transient states:

$$\phi_D : \square_{<} \neg r \wedge \diamond_{\leq \epsilon} r \wedge \square_{>0} (r \rightarrow (\neg r) \mathcal{U}_{=\delta} r).$$

Let ϕ^* be the MITL-formula that results from ϕ by replacing every occurrence of a subformula $\psi_1 \mathcal{U}_I \psi_2$ with

$$\neg r \wedge (\psi_1 \wedge \neg r) \mathcal{U}_{\geq 0} \psi_2$$

if I is $[0, 0]$; with

$$(r \wedge (\psi_1 \wedge \neg r) \mathcal{U}_{\geq 0} \psi_2) \vee \psi_1 \mathcal{U}_{(0,\delta)} (r \wedge \psi_1 \wedge \psi_1 \mathcal{U}_{(0,\delta]} \psi_2)$$

if I is $[1, 1]$; with

$$\psi_1 \mathcal{U}_{\{(I-1)\delta, \tau(I)\delta\}} (r \wedge \psi_1 \wedge \psi_1 \mathcal{U}_{(0,\delta]} \psi_2)$$

if I is bounded and $Z(I) > 1$; and with

$$\psi_1 \mathcal{U}_{\geq (I-1)\delta} (r \wedge \psi_1 \wedge \psi_1 \mathcal{U} \psi_2)$$

if I is unbounded and $Z(I) > 1$. It is not hard to show that $\rho_D \models \phi$ iff $\rho \models \phi_D \wedge \phi^*$ for every timed state sequence ρ .

For example, consider the MTL-formula

$$\square_{\geq 0} (p \rightarrow \diamond_{=5} q),$$

and the digital clock $D = (1, 0)$. This formula D -specifies the property that “for every p -state there is a q -state separated from p by exactly five integer times,” and is equivalent to the MITL-formula

$$\phi_{(1,0)} \wedge \square_{\geq 0} (p \rightarrow \diamond_{[4,5]} (r \wedge \diamond_{(0,1]} q)).$$

(b) From the tableau decision procedure for MTL [AH90], it follows that if a formula ϕ of MTL is satisfiable, then it has a model ρ_D such that any two state changes in ρ are separated by at least some minimum time gap (which depends on D and the size of ϕ). In fact, for any digital clock D one can always construct timed state sequences in Π_ϕ^D that become periodic after some point in time. We show that this is not the case for MITL (although, as we shall see later, it is the case that any satisfiable MITL-formula has a model in which in any fixed interval of time there is only a bounded number of state changes).

Let us construct a satisfiable MITL-formula ϕ with the property that every model $\rho = (\sigma, \tau)$ of ϕ contains arbitrarily close state changes; that is, for every real $\delta > 0$, there is some $i \geq 1$ such that $s_{i-1} \neq s_i$ and $s_i \neq s_{i+1}$ and $t_{i+1} - t_i < \delta$. The set of models of ϕ can clearly not be specified in MTL, for any choice of digital clock D .

The formula ϕ uses three propositions p , q , and r . First, it requires at most one of these three propositions to be true at any state. In addition, it has the following three conjuncts. The first condition,

$$r \wedge \square \diamond \square \square \square \rightarrow (\neg r) \mathcal{U}_{=2} r,$$

places transient r -states at precisely the even integers. The second condition,

$$\square_{\geq 0} ((p \vee q) \rightarrow \diamond_{<1} r),$$

ensures that p and q can only hold in the second half of the intervals of length 2 separating consecutive r -states. The third condition,

$$\diamond_{<2} p \wedge \square_{\geq 0} (p \rightarrow \diamond_{<1} q) \wedge \square_{\geq 0} (q \rightarrow \diamond_{(2,3)} p),$$

implies that there is a p -state, and later a q -state, between every pair of consecutive r -states, and thus between every odd integer and the subsequent even integer.

Moreover, from any model of ϕ we can extract an infinite sequence of alternating p and q states, with the q -state following a p -state guaranteed by the condition $p \rightarrow \diamond_{<1} q$, and the p -state following a q -state by the condition $q \rightarrow \diamond_{(2,3)} p$. The times that are associated with the states in this sequence, taken modulo 2, form a strictly increasing infinite sequence of reals contained in the interval $(1, 2)$. Since this time sequence is bounded above, there must be arbitrarily close pairs of a p -state followed by a q -state. It follows that ϕ has no eventually periodic models.

On the other hand, the MITL-formula ϕ is satisfiable; a model for ϕ can be readily constructed by introducing, in addition to the transient r-states at all even integers, transient p-states at time $2n - 2/4^n$, and transient q-states at time $2n - 1/4^n$, for each integer $n \geq 1$. ■

3 Timed Automata

We use a variant of timed automata defined in [AD90] to model finite-state real-time systems. This formalism is a generalization of (nondeterministic) finite-state machines over infinite strings. While w-automata generate (or accept) infinite sequences of states [Tho90], timed automata are additionally constrained by timing requirements and produce *timed* state sequences.

A timed automaton operates with finite control — a finite set of states and a finite set of real-valued clocks. All clocks proceed at the same rate and measure the amount of time that has elapsed since they were started (or reset). Each transition of the automaton may reset some of the clocks; each state of the automaton puts certain constraints on the values of the atomic propositions as well as on the values of the clocks: the control of the automaton can reside in a particular state only if the values of the propositions and clocks satisfy the corresponding constraints.

We permit only simple constraints on the clock values. A *clock constraint* $\mathcal{I} \subseteq \mathbb{R}^+$ is a finite union of (possibly unbounded) intervals with rational endpoints; the value $y(z) \in \mathbb{R}^+$ of a clock x satisfies the constraint \mathcal{I} iff $y(c) \in \mathcal{I}$. We usually denote the clock constraints for a clock x as boolean combination of arithmetic expressions containing x ; for instance,

$$1 \leq x < 3 \vee x = 4 \vee x > 5$$

stands for the clock constraint $[1, 3) \cup [4, 4] \cup (5, \infty)$ that restricts the value of x . Let \mathcal{R} be the set of clock constraints.

Formally, a *timed automaton* is a six-tuple $M = \langle S, C, \mu, \nu, S_0, E \rangle$, where

S is a finite set of states,

C is a finite set of clocks,

$\mu: S \rightarrow 2^P$ assigns to each state and proposition a truth value,

$\nu: S \rightarrow \mathcal{R}^C$ assigns to each state and clock a clock constraint,

$S_0 \subseteq S$ is a set of initial states,

$E \subseteq S^2 \times 2^C$ is a set of transitions. Each transition $\langle s, s', \lambda \rangle$ identifies a source state s , a target state s' , and a set $\lambda \subseteq C$

of clocks to be reset; we usually denote this transition by $s \xrightarrow{\lambda} s'$.

The runs of a timed automaton define timed state sequences. At any time instant during a run, the configuration of the automaton is completely determined by the state in which the control resides and the values of all clocks. The values of all clocks are given by a *clock interpretation* γ , which is a map from C to \mathbb{R}^+ : for any clock $x \in C$, the value of x under the interpretation γ is $\gamma(x) \in \mathbb{R}^+$.

Assume that, at time $t \in \mathbb{R}^+$, a timed automaton is in state s and the clock values are given by the clock interpretation γ . Suppose that the state of the automaton remains unchanged during the time interval I with $Z(I) = t$. All clocks proceed at the same rate as time elapses; at any time $t' \in I$ the value of any clock x is $\gamma(x) + t' - t$. During all this time the value of x satisfies the clock constraint that is associated with s and x :

$$(\gamma(x) \upharpoonright t' - t) \in \nu(s, x).$$

Now suppose that the automaton changes its state at time $r(1) = t''$ via the transition $s \xrightarrow{\lambda} s'$. This state change happens in one of two ways. If I is right-closed, then the state at time t'' is still s and

$$(\gamma(x) \upharpoonright t'' - t) \in \nu(s, x)$$

for all clocks x ; otherwise the state at time t'' is s' and $0 \in \nu(s', x)$ for all clocks $x \in \lambda$, which are reset, and

$$(\gamma(x) + t'' - t) \in \nu(s', x)$$

for all other clocks.

Let us formalize this intuition. Suppose we are given a timed automaton $M = (S, C, \mu, \nu, S_0, E)$; a *run* of M is a finite or infinite sequence

$$r: \quad \xrightarrow[\gamma_0]{} (s_0, I_0) \xrightarrow[\gamma_1]{\lambda_1} (s_1, I_1) \xrightarrow[\gamma_2]{\lambda_2} (s_2, I_2) \xrightarrow[\gamma_3]{\lambda_3} \dots$$

of states $s_i \in S$, intervals I_i , clock sets $\lambda_i \subseteq C$, and clock interpretations $\gamma_i : C \rightarrow \mathbb{R}^+$ such that

- . $s_0 \in S_0$,
- . $(s_i, s_{i+1}, \lambda_i) \in E$ for all $i \geq 0$,
- . $I_0 I_1 I_2 \dots$ is an interval sequence,
- . for all $x \in C$ and $i \geq 0$, we have $\gamma_{i+1}(x) = 0$ if $x \in \lambda_{i+1}$, and $\gamma_{i+1}(x) = \gamma_i(x) + r(I_i) - l(I_i)$ otherwise.
- . $(\gamma_i(x) \upharpoonright t - l(I_i)) \in \nu(s_i, x)$ for all $x \in C$, $i \geq 0$, and $t \in I_i$.




Figure 1: Timed automaton

Note that, according to this definition, the clocks may start at any real values that satisfy the clock constraints of an initial state.

The run r uniquely determines the timed state sequence

$$\rho_r: (\mu(s_0), I_0) \rightarrow (\mu(s_1), I_1) \rightarrow (\mu(s_2), I_2) \rightarrow \dots$$

By $\Pi(\mathcal{M})$ we denote the set of all timed state sequences ρ_r that correspond to runs of the timed automaton M . We say that M *generates* (or *accepts*) the timed state sequences in $\Pi(\mathcal{M})$.

We will use timed automata to model real-time systems. A real-time system is represented by the timed automaton M iff its possible behaviors are exactly the timed state sequences in $\Pi(M)$. Accordingly, the system modeled by M satisfies its MITL-specification ϕ , denoted by $M \models \phi$, iff $\rho_r \models \phi$ for all runs r of M .

We point out that a run may contain transient states. Such states allow us to model instantaneous conditions during the execution of a real-time system, like the occurrence of events. Their times can be enforced accurately by using singular intervals as clock constraints.

Consider, for example, the timed automaton M in Figure 1. The automaton M has six states, s_0 to s_5 , and uses two clocks, x and y . The label $x := 0$ on a transition indicates that the clock x is reset by that transition.

The automaton starts in the initial state s_0 with the clock y initialized to 0. At time 40 the automaton moves to state s_5 , and simply loops there. The

proposition p denotes an external event which is true only at instantaneous points $t < 40$ in time (and no more than once every 5 time units), namely, whenever M is in state s_2 . The automaton responds to p by resetting the clock x , and then it requires that the proposition q holds over the interval $t + [2, 5)$. Thus the automaton M models a system which responds, until time 40, to the event p by setting q to true for the interval $[2, 5)$ following p . A possible timed state sequence generated by M is

$$(\emptyset, [0, 13)) \rightarrow (\{p\}, [13, 13]) \rightarrow (\emptyset, (13, 15)) \rightarrow (\{q\}, [15, 20)) \rightarrow (\emptyset, [20, 40)) \rightarrow (\{q\}, [40, \infty)).$$

The emptiness problem for timed automata is solved in [AD90]: the problem of whether a timed automaton has any run is PSPACE-complete. Our definition of timed automata is somewhat more general than the one in [AD90]; it can also enforce transient states. But the decision procedure for checking emptiness can be easily adapted to prove the following result:

Theorem 3.1 (Emptiness of timed automata) *The problem of deciding if $\Pi(\mathcal{M}) = \emptyset$ for a timed automaton $M = (S, C, \mu, \nu, S_0, E)$ is PSPACE-complete. Moreover, there is an algorithm that decides this problem in time $O((|S| + |E|) \cdot 2^{|\nu|})$.*

To enforce *fairness* constraints on the legal behaviors of a real-time system, we add standard liveness conditions to timed automata, such as *Büchi* acceptance criteria or *Muller* acceptance criteria for w-automata (see [AD90] for details). Theorem 3.1 carries over to either case.

4 Deciding MITL

We solve the satisfiability problem for MITL by reducing it to the emptiness problem for timed automata. Our main result is that, given an MITL-formula ϕ , we can construct a timed automaton \mathcal{M}_ϕ such that the runs of \mathcal{M}_ϕ that meet certain fairness requirements correspond precisely to the timed state sequences that satisfy ϕ .

4.1 Restricting the problem

To simplify the exposition of the decision procedure, we restrict the satisfiability question for MITL to formulas and models of a specific form and show that this can be done without loss of generality.

Given an MITL-formula ϕ , a timed state sequence ρ , and a constant $a \in \mathbb{Q}$, let $a\phi$ and $a\rho$ be the MITL-formula and the timed state sequence that result from ϕ and ρ , respectively, by replacing each interval I by the interval aI . Clearly, $\rho \models \phi$ iff $a\rho \models a\phi$. Thus, for the purpose of checking the satisfiability of ϕ , we may assume that all interval end-points in ϕ are integers; for if they are not, then consider $n_\phi\phi$ for the least common denominator n_ϕ of all (rational) interval end-points in ϕ . This translation causes at most a quadratic blow-up in the size of the formula.

Next we give a series of transformations that allow us to rewrite any formula ϕ into an equivalent formula ϕ^* that contains only temporal operators of very specific forms.

First, we require that no interval in ϕ contains 0. This can be achieved by applying the following equivalence:

$$\psi_1 \mathcal{U}_I \psi_2 \leftrightarrow (\psi_2 \vee \psi_1 \mathcal{U}_{I \cap (0, \infty)} \psi_2)$$

provided that $0 \in I$.

Secondly, we require that the only unbounded intervals in ϕ are of the form $(0, \infty)$. This can be achieved by applying the following two equivalences:

$$\psi_1 \mathcal{U}_{(n, \infty)} \psi_2 \leftrightarrow \Box_{(0, n]} (\psi_1 \wedge \psi_1 \mathcal{U} \psi_2)$$

$$\psi_1 \mathcal{U}_{[n, \infty)} \psi_2 \leftrightarrow \Box_{(0, n)} \psi_1 \wedge \Box_{(0, n]} (\psi_2 \vee (\psi_1 \wedge \psi_1 \mathcal{U} \psi_2))$$

provided that $n > 0$.

Thirdly, we require that only the *eventually* and the *always* operators are constrained with bounded intervals I such that $Z(1) = 0$. This can be achieved by applying the following equivalence:

$$\psi_1 \mathcal{U}_I \psi_2 \leftrightarrow \Diamond_I \psi_2 \wedge \psi_1 \mathcal{U} \psi_2$$

provided that $Z(1) = 0$.

Finally, we push all negations in ϕ to the inside and use the following equivalence to eliminate each subformula of the form $\psi_1 \cup \psi_2$:

$$\psi_1 \cup \psi_2 \leftrightarrow \Box \psi_1 \vee \psi_1 \mathcal{U} (\psi_1 \wedge \psi_2).$$

The resulting formula ϕ^* is equivalent to ϕ and consists of atomic propositions, negated atomic propositions, conjunctions, disjunctions, and temporal subformulas ψ of the following six types:

1. $\psi_1 \mathcal{U}_I \psi_2$ with bounded I and $l(I) > 0$.
2. $\psi_1 \text{ }_I\mathcal{U} \psi_2$ with bounded I and $Z(1) > 0$.
3. $\diamond_I \psi'$ with $I = (0, n)$ or $I = (0, n]$.
4. $\square \text{ }_I \psi'$ with $I = (0, n)$ or $I = (0, n]$.
5. $\psi_1 \mathcal{U} \psi_2$.
6. $\square \psi'$.

Although these rewritings blow up the size of the formula ϕ , we can bound the size of the constants in ϕ^* and the number of subformulas in ϕ^* as follows:

- Let $K \in \mathbb{N}$ be such that $K - 1$ is the largest (integer) constant appearing as an interval end-point in ϕ . Then the largest constant that occurs as an end-point of an interval in ϕ^* is $K - 1$.
- Let $N \in \mathbb{N}$ be the number of atomic propositions, boolean connectives, and temporal operators in ϕ . Then the number of syntactic subformulas of ϕ^* is $O(N)$.

Thus we restrict ourselves to test the satisfiability of MITL-formulas each of whose temporal subformulas are, according to the above classification, of one of six types, *type-1* to *type-6*.

Moreover, to check the satisfiability of an MITL-formula ϕ , by Lemma 2.1 we can confine ourselves to the question if ϕ has a $\&$ fine model. Therefore we consider, throughout this section, only $\&$ fine timed state sequences $\rho = (\sigma, \tau)$. It follows that, if ψ is a subformula of ϕ , we may write $\rho^i \models \psi$ for “ $\rho^t \models \psi$ for all $t \in I_i$.” In addition, we assume that all intervals in τ are either singular or open. This is sufficient, because any model of ϕ can be brought into this form by splitting all nonsingular (half)closed intervals; for instance, the interval $[a, b)$ can be split into the two intervals $[a, a]$ and (a, b) .

Let us introduce a new atomic proposition p_{sing} such that $\rho^i \models p_{sing}$ iff the i -th interval I_i of $\rho = (a, \tau)$ is singular. Hence the proposition p_{sing} holds exactly in every other interval. For a timed state sequence ρ that satisfies these conditions and $t \in \mathbb{R}^+$, let i be such that $t \in I_i$. Then:

$\rho^t \models_{t=} \psi_1 \mathcal{U}_I \psi_2$ iff $\rho^i \models \psi_1 \vee p_{sing}$, and both $\rho^j \models \psi_2$ and $\rho^j \models \psi_1 \vee p_{sing}$ for some j with $I_j \cap (t + I) \neq \emptyset$, and $\rho^k \models \psi_1$ for all $i < k < j$.

$\rho^t \models \psi_1 \text{ I U } \psi_2$ iff $\rho^i \models \psi_1$ if $I_i \cap (t + I) \neq \emptyset$, and either $\rho^i \models \psi_2 \wedge \neg p_{sing}$, or $\rho^j \models \psi_2$ for some $j > i$ and $\rho^k \models \psi_1$ for all $i < k \leq j$ with $I_k \cap (t + I) \neq \emptyset$, or $\rho^k \models \psi_1$ for all $k > i$ with $I_k \cap (t + I) \neq \emptyset$.

The different types of temporal subformulas of ϕ are handled differently by our algorithm. The simplest case is that of type-5 and type-6 formulas; they are treated essentially in the same way in which tableau decision procedures for linear temporal logic handle unconstrained temporal operators. The most interesting case is that of type-1 and type-2 formulas. We concentrate first on this case. The case of type-3 and type-4 formulas will be considered later.

4.2 Outline of the algorithm

Consider the MITL-formula

$$\square_{[0,1)} (p \rightarrow \diamond_{[1,2]} q).$$

Let us assume that both p and q are true only in singular intervals and let us try to build a timed automaton that accepts precisely the models of this formula.

Whenever the automaton visits a p -state, it needs to make sure that within 1 to 2 time units a q -state is visited. This can be done by setting a clock x to 0 when the p -state is visited, and demanding that some q -state with the clock constraint $1 \leq x \leq 2$ is visited later. This strategy requires a clock per visit to a p -state within the interval $[0, 1)$. However, the number of such visits is potentially unbounded and, hence, any automaton with a fixed number of clocks cannot reset a new clock for every visit. That is why this simple strategy cannot be made to work.

An alternative approach is to guess the times for future q -states in advance. The automaton nondeterministically guesses two time values t_1 and t_2 within the interval $[0, 1)$; this is done by resetting a clock x at time t_1 and another clock y at time t_2 . The guess is that the *last* q -state within the interval $[1, 2)$ is at time $t_1 + 1$, and that the *first* q -state within the interval $[2, 3)$ is at time $t_2 + 2$. If the guesses are correct, then the formula $\diamond_{[1,2]} q$ holds during the intervals $[0, t_1]$ and $[t_2, 1)$, and does not hold during the

interval (t_1, t_2) . Consequently, the automaton requires that every p-state within the interval $[0, 1)$ lies either within $[0, t_1]$ or within $[t_2, 1)$. It also needs to make sure that the guesses are right; that is, whenever either $x = 1$ or $y = 2$, the automaton must be in a q-state. This strategy requires only two clocks for the interval $[0, 1)$ of length 1, irrespective of the number of p-states within $[0, 1)$.

We say that the guessed times $t_1 + 1$ and $t_2 + 2$ witness the formula $\diamond_{[1,2]} q$ throughout the intervals $[0, t_1]$ and $[t_2, 1)$, respectively. In general, the witnesses need not be singular intervals, they can be open intervals. In the following we develop an algorithm based on this idea of guessing, in advance, time intervals that witness temporal formulas and, later, checking the correctness of these guesses. The crucial fact that makes this strategy work, with a finite number of clocks, is that the *same* interval may serve as a witness for many points in time.

4.3 Witnessing intervals

The interval I' is called a *witnessing interval* for the MITL-formula $\psi_1 \mathcal{U}_I \psi_2$ under ρ^t , for a timed state sequence ρ and $t \in \mathbb{R}^+$, iff $I' \cap (t + I) \neq \emptyset$ and $\rho^t \models \psi_1 \mathcal{U}_{J-t} \psi_2$ for every nonempty interval $J \subseteq I'$. Observe that if I' witnesses $\psi_1 \mathcal{U}_I \psi_2$ under ρ^t , then $\rho^{t'} \models \psi_1$ for all $t < t' < r(I')$ and $\rho^{t'} \models \psi_2$ for all $t' \in I'$. The interval I' is a witnessing interval for the MITL-formula $\psi_1 I \cup \psi_2$ under ρ^t iff $t + I \subseteq I'$ and $\rho^t \models \psi_1 I'_{-t} \cup \psi_2$.

Witnessing intervals are defined such that the following property holds:

Lemma 4.1 (Witnessing intervals) *Let ψ be an MITL-formula of the form $\psi_1 \mathcal{U}_I \psi_2$ or $\psi_1 I \cup \psi_2$, let p be a timed state sequence and $t \in \mathbb{R}^+$. There is a witnessing interval for ψ under ρ^t iff $\rho^t \models \psi$.*

Proof of Lemma 4.1 If $\rho^t \models \psi$ for the formula $\psi = \psi_1 \mathcal{U}_I \psi_2$, then $\rho^{t'} \models \psi_2$ for some $t' \in t + I$ and the singular interval $[t', t']$ witnesses ψ under ρ^t . If $\rho^t \models \psi$ for the formula $\psi = \psi_1 I \cup \psi_2$, then the interval $t + I$ witnesses ψ under ρ^t .

The other direction of the lemma follows from the semantic clauses for the *until* and *unless* operators. ■

Now we show that the same interval may serve as a witnessing interval for a temporal formula under (infinitely) many suffixes of a timed state sequence.

Consider, for example, the timed state sequence ρ over two propositions p and q :

$$(\{p\}, [0, 1.2]) \rightarrow (\{p, q\}, (1.2, 1.6)) \rightarrow (\{p\}, [1.6, \infty))$$

Thus along ρ the proposition p is always true, but the proposition q is true only during the interval $I_q = (1.2, 1.6)$. The interval I_q witnesses the formula $p\mathcal{U}_{(1,2)}q$ under ρ^t for every $t \in [0, 0.6)$. On the other hand, the interval $[1.6, 3]$ witnesses the formula $\Box_{(1,2)}(\neg q)$ under ρ^t for every $t \in [0.6, 1]$.

Lemma 4.2 (Sharing type-1 witnesses) *Let ψ be the type-1 MITL-formula $\psi_1\mathcal{U}_I\psi_2$. For every timed state sequence ρ , there are two bounded intervals I_1 and I_2 such that, for every $t \in [0, 1)$, the formula ψ is satisfied by ρ^t iff either I_1 or I_2 witnesses ψ under ρ^t . Furthermore, I_i is either singular or open, and $r(I_i) \leq r(I) + 1$ for $i = 1, 2$.*

Proof of Lemma 4.2 Let $\rho = (a, \tau)$ be a ψ -fine timed state sequence with only singular and open intervals, including the singular interval $[r(I) + 1, r(I) + 1]$ (split intervals if necessary). We choose two witnessing intervals I_1 and I_2 as follows:

- Let \hat{i} be the *maximal* $i \geq 0$ such that $I_i \cap I \neq \emptyset$, both $\rho^i \models \psi_2$ and $\rho^i \models \psi_1 \vee p_{sing}$, and $\rho^k \models \psi_1$ for all $0 \leq k < i$ with $I_k \cap I \neq \emptyset$. If no such i exists, let $I_1 = \emptyset$; otherwise, let $I_1 = I_{\hat{i}}$.
- Let \hat{j} be the *minimal* $j \geq 0$ such that $I_j \cap (I + 1) \neq \emptyset$, both $\rho^j \models \psi_2$ and $\rho^j \models \psi_1 \vee p_{sing}$, and $\rho^k \models \psi_1$ for all $0 \leq k < j$ with $I_k \cap (I \cup I + 1) \neq \emptyset$. If no such j exists, let $I_2 = \emptyset$; otherwise, let $I_2 = I_{\hat{j}}$.

Assume that $0 \leq t < 1$; then ρ^t satisfies ψ iff $\rho^{t'} \models \psi_1$ for all $t < t' < I$ and either $I_{\hat{i}} \cap (t + 1) \neq \emptyset$ or $I_{\hat{j}} \cap (t + I) \neq \emptyset$. The first case is equivalent to I_1 witnessing ψ under ρ^t ; the second case is equivalent to I_2 witnessing ψ under ρ^t . ■

In the case of type-2 formulas, a single witness per unit interval suffices to reduce the problem to type 3:

Lemma 4.3 (Sharing type-2 witnesses) *Let ψ be the type-2 MITL-formula $\psi_1 I \cup \psi_2$. For every timed state sequence ρ , there is a bounded interval I' such that, for every $t \in [0, 1)$, the formula ψ is satisfied by ρ^t iff either ρ^t satisfies the type-3 formula $\Diamond_{(0,\infty) \cap (<I)} \psi_2$ or I' witnesses ψ under ρ^t . Furthermore, $r(I') \leq r(I) + 1$.*

Proof of Lemma 4.3 Let $\rho = (\sigma, \tau)$ be a ψ -fine timed state sequence with only singular and open intervals, including the singular interval $I_n = [r(I) \$1, r(I) + 1]$ (split intervals if necessary). We choose witnessing interval I' as follows:

- Let \hat{i} be the minimal $i \geq 0$ such that $I_i \cap I \neq \emptyset$ and either
 1. $\rho^k \models \psi_1$ for all $k \geq i$ with $I_k \cap I \neq \emptyset$, or
 2. there is some $i \leq j \leq n$ such that $\rho^j \models \psi_1 \wedge \psi_2$ and $\rho^k \models \psi_1$ for all $i \leq k < j$.
- Given \hat{i} , let \hat{j} be the maximal $\hat{i} \leq j \leq n$ such that either $\rho^k \models \psi_1$ for all $\hat{i} \leq k \leq j$, or $\rho^k \models \psi_1 \wedge \psi_2$ for some $\hat{i} \leq k \leq j$. Note that if \hat{i} exists, then so does \hat{j} ; in particular, if \hat{i} exists because of clause 2, then $\hat{j} = n$.

If no appropriate \hat{i} exists, let $I' = \emptyset$; otherwise, let I' be the union of all I_k for $\hat{i} \leq k \leq \hat{j}$.

Assume that $0 \leq t < 1$; then ρ^t satisfies ψ iff either (1) $\rho^i \models \psi_1$ for all i with $I_i \cap (t + I) \neq \emptyset$, or (2) $\rho^i \models \psi_1 \wedge \psi_2$ for some i with $I_i \cap (t + I) \neq \emptyset$ and $\rho^j \models \psi_1$ for all $j < i$ with $I_j \cap (t + I) \neq \emptyset$, or (3) $\rho^{t'} \models \psi_2$ for some $t < t' < t + I$. In either of the first two cases, I' witnesses ψ under ρ^t ; the third case is equivalent to ρ^t satisfying the formula $\diamond_{(0, \infty) \cap (< I)} \psi_2$. If I' witnesses ψ under ρ^t , then $\rho^t \models \psi$ by Lemma 4.1. ■

4.4 Type-1 and type-2 formulas

Now we can be more precise about how we will construct the timed automaton \mathcal{M}_ϕ that accepts exactly the models of ϕ . To check the truth of type-1 and type-2 subformulas of ϕ , the automaton guesses corresponding witnessing intervals. The boundaries of a witnessing interval are marked by clocks: a *clock interval* is a bounded interval that is defined by its *type* (e.g., left-closed and right-open) and a pair of clocks. Given a time t and a clock interpretation γ , the clock interval $C = [x, y]$, for two clocks x and y , stands for the closed witnessing interval $[t + K - \gamma(x), t + K - \gamma(y)]$; the clock interval $C = [x, y)$ stands for the corresponding half-open interval, etc. We write $K - C$ for the interval $\{K - \gamma(z), K - \gamma(y)\}$, for any type of clock interval $C = \{x, y\}$.

For simplicity, let us consider a type-1 subformula ψ of the form $\diamond_I \psi'$. The automaton resets, nondeterministically, any of its clocks at any time. When guessing a witnessing interval I' , it writes the prediction that “the

clock interval $C = \{x, y\}$ witnesses the formula ψ into its memory. If the clock x was reset at time t_1 , and y was reset at time $t_2 \geq t_1$, then the witnessing interval guessed is $I' = \{t_1 + K, t_2 + K\}$. To check the truth of the temporal formula ψ at time $t \geq t_2$, the automaton needs to verify that its guess I' is indeed a witness. The condition $I' \cap (t + 1) \neq \emptyset$ translates to verifying the clock constraint $(K - C) \cap I \neq \emptyset$. It remains to be checked that ψ' is satisfied throughout the witnessing interval I' ; that is, the automaton needs to verify that ψ' holds at all states with the clock constraint $0 \in (K - C)$.

The Lemmas 4.2 and 4.3 are the key to constructing an automaton that needs only *finitely* many clocks. For the type-1 formula $\psi_1 \mathcal{U}_I \psi_2$, at most two witnessing intervals need to be guessed per interval of unit length. Furthermore, the fact that the right end-point of a witnessing interval is bounded allows the automaton to reuse every clock after a period of length $r(I) + 1$. Thus we need, at any point in time, at most $2r(I) + 2$ *active* clock intervals; that is, clock intervals that stand for a guess of a witnessing interval and, therefore, have to be verified later. Similarly, to check a type-2 formula $\psi_1 \mathcal{I} \psi_2$, we need, at any point in time, no more than $r(I) + 1$ active clock intervals. Consequently, $2K$ clock pairs suffice to check any type-1 subformula of ϕ , and K clock pairs suffice for any type-2 subformula of ϕ .

4.5 Type-3 and type-4 formulas

Now let us move to formulas of the form $\diamond_I \psi'$ and $\square_I \psi'$ with $I = (0, n)$ or $I = (0, n]$. Checking the truth of such a formula is much easier and can be done using a single clock.

Consider the type-3 formula $\psi = \diamond_I \psi'$. Whenever the automaton needs to check that ψ holds, say at time t , it starts a clock x and writes the corresponding proof obligation into its memory — to verify that ψ' holds at some later state with the clock constraint $x \in I$. The obligation is discharged as soon as an appropriate ψ' -state is found. If the automaton encounters another ψ -state in the meantime, at time $t' > t$ before the obligation is discharged, it does not need to check the truth of ψ separately for this state. This is because if there is a ψ' -state after time t' within the interval $t + I$, then both $\rho^t \models \diamond_I \psi'$ and $\rho^{t'} \models \diamond_I \psi'$. Once the proof obligation is discharged, the clock x can be used again. Thus one clock suffices to check the formula ψ as often as necessary.

The described strategy works for checking the truth of ψ at singular intervals. There is, however, a subtle problem with this method when the

truth of ψ during open intervals needs to be checked, as is illustrated by the following example. Consider the timed state sequence

$$(\{\}, [0, 0]) \rightarrow (\{\}, (0, 1)) \rightarrow (\{p\}, [1, \infty]);$$

it satisfies the formula $\diamond_{(0,1)}p$ at all times $t \in (0, 1)$. To check the truth of $\diamond_{(0,1)}p$ during the open interval $(0, 1)$, the automaton starts a clock x upon entry, at time 0. However, the proof obligation that p holds at some later state with the clock constraint $x \in I$ can never be verified. On the other hand, if the automaton were to check, instead, the truth of the formula $\diamond_{(0,1]}p$ during the interval $(1, 0)$, then our strategy works and the corresponding proof obligation can be verified, because there is a p -state while $x \in (0, 1]$ holds. Furthermore, observe that the validity of $\diamond_{(0,1]}p$ throughout the open interval $(0, 1)$ implies that $\diamond_{(0,1)}p$ is also true throughout $(0, 1)$.

In general, the following lemma holds:

Lemma 4.4 (Weakening type-3 formulas) *Let ψ and $\hat{\psi}$ be the type-3 MITL-formulas $\diamond_I \psi'$ and $\diamond_{I \cup \{r(I)\}} \psi'$, respectively. For every timed state sequence $\rho = (\sigma, \tau)$ and open interval I_i in τ , $\rho^i \models \psi$ iff $\rho^i \models \hat{\psi}$.*

Proof of Lemma 4.4 First note that, for all $t \geq 0$, if ψ is satisfied by ρ^t , then $\hat{\psi}$ is also satisfied by ρ^t . This is because $I \subseteq I \cup \{r(I)\}$.

Now consider an open interval I_i and assume that $\rho^i \models \hat{\psi}$. If I is right-closed, then $\psi = \hat{\psi}$. So suppose that I is right-open, and let $t \in I_i$. Since I_i is open, there exists some $t' \in I_i$ with $t' < t$. Since $\rho^{t'} \models \hat{\psi}$, there exists some $j \geq i$ such that $I_j \cap (t' + (I \cup \{r(I)\})) \neq \emptyset$ and $\rho^j \models \psi'$. It follows that $I_j \cap (t + I) \neq \emptyset$ and, hence, that $\rho^t \models \psi$. ■

Consequently, to check the truth of a type-3 formula ψ during an open interval, it suffices to check the truth of the weaker formula $\hat{\psi}$. Accordingly, the automaton we construct writes only the proof obligation that corresponds to checking $\hat{\psi}$ into its memory.

For checking a type-4 formula of the form $\psi = \square I+$, the situation is symmetric. The automaton uses also a single clock x to check this formula. Whenever the formula ψ needs to be verified, say at time t , the automaton starts the clock x with the proof obligation that as long as the clock constraint $x \in I$ holds, so does ψ' . The obligation is discharged as soon as $x > I$. If the automaton encounters another ψ -state within the interval $t + I$, say at time t' , it simply resets the clock x , and thus overwrites the previous proof obligation. This strategy is justified by the observation that if ψ'

holds throughout the interval $(t, t']$ and $\rho^{t'} \models \Box \psi'$, then also $\rho^t \models \Box \psi$. Once the proof obligation is discharged, the clock x can be reused to check ψ again whenever necessary.

As in the case of type-3 formulas, we need to be more careful when checking ψ during open intervals. For the type-4 formula $\psi = \Box \text{IS}'$, let $\hat{\psi}$ be the formula $\Box \text{I}\{-\text{I}\} \psi'$. From Lemma 4.4 and duality, it follows that for every timed state sequence $p = (\sigma, \tau)$, if I_i is open, then $\rho^i \models \psi$ iff $\rho^i \models \hat{\psi}$. Hence to check the truth of ψ during an open interval, it suffices again to check the truth of the weaker formula $\hat{\psi}$. Accordingly, only a proof obligation for $\hat{\psi}$ is set up. This is because the corresponding clock x is started at time $r(I_i)$, and for ψ to hold during the open interval I_i , ψ' need not hold at time $r(I_i) + r(I)$, even if I is right-closed.

4.6 Constructing the timed automaton

Now let us define the timed automaton \mathcal{M}_ϕ formally. For each temporal subformula of ϕ of type-1, the automaton \mathcal{M}_ϕ has $2K$ pairs of clocks. These clocks always appear in pairs, to form clock intervals. From any pair of clocks x and y , four different clock intervals can be formed: (x, y) , $[x, y)$, $(x, y]$, and $[x, y]$. According to Lemma 4.2, for checking type-1 formulas we need only singular and open witnessing intervals. Thus associated with each type-1 subformula ψ of ϕ we have $4K$ clock intervals; they are denoted by $C_1(\psi), \dots, C_{4K}(\psi)$. For each type-2 subformula of ϕ the automaton uses K clock pairs giving $4K$ clock intervals. For subformulas ψ of types 3 and 4, the automaton needs one clock x_ψ per formula.

In addition to these clocks, we use the clock x_{sing} to enforce that the runs of \mathcal{M}_ϕ have alternate singular and open intervals.

Given the MITL-formula ϕ , we define its *closure set* $Closure(\phi)$ to consist of the following items:

1. All subformulas of ϕ .
2. For each type-2 formula $\psi_1 \text{I}\cup \psi_2$ in the closure set, the type-3 formula $\Diamond_{(0, \infty) \cap (< I)} \psi_2$; for each type-3 formula $\psi = \Diamond_I \psi'$ in the closure set, the type-3 formula $\hat{\psi} = \Diamond_{I \cup \{r(I)\}} \psi'$; and for each type-4 formula $\psi = \Box \text{IS}'$ in the closure set, the type-4 formula $\hat{\psi} = \Box_{I - \{r(I)\}} \psi'$.
3. For each type-1 formula ψ in the closure set, the clock intervals $C_1(\psi), \dots, C_{4K}(\psi)$; for each type-2 formula ψ in the closure set, the clock

intervals $C_1(\psi), \dots, C_{4K}(\psi)$; and for each type-3 and type-4 formula ψ in the closure set, the clock x_ψ .

4. For each clock interval $C = C_j(\psi)$ in the closure set, where ψ is $\psi_1 \mathcal{U}_I \psi_2$ or $\psi_1 \mathcal{I} \psi_2$, all clock constraints of the form $0 < (K - C)$, $0 \subset (K - C)$, $0 = (K - C)$, $(K - C) = \emptyset$, $I \subseteq (K - C)$, and $(K - C) \cap I \neq \emptyset$; and for each clock x_ψ in the closure set, where ψ is $\diamond_I \psi'$ or $\square_I \psi'$, the clock constraints $x \in I$ and $x > I$.

We write $0 \subset (K - C)$ short for $\{0\} \subset (K - C)$. It should be clear that all of these conditions are indeed clock constraints. For instance, the condition $0 \subset (K - [x, y))$ stands for the clock constraint $x \leq K \wedge y > K$; the condition $0 = (K - [x, y))$ is never satisfied.

5. The clock constraint $x_{sing} = 0$.

Note that the number of subformulas of ϕ is $O(N)$ and the number of clocks is $0(K)$ for each subformula of ϕ . Hence the size of the closure set $Closure(\phi)$ is $O(N \cdot K)$.

The states of the desired automaton \mathcal{M}_ϕ will be subsets of $Closure(\phi)$. We need to consider only those subsets of $Closure(\phi)$ that satisfy certain local consistency constraints. Whenever the automaton is in state s , the formulas in s indicate which subformulas of ϕ are true. Accordingly, a state $s \subseteq Closure(\phi)$ is initial iff both ϕ and $x_{sing} = 0$ are in s , and for each state s the propositional constraints $\mu(s)$ are defined such that $p \in \mu(s)$ iff $p \in s$ for all atomic propositions $p \in P$.

The clock constraints $\nu(s)$ are the conjunction of all clock constraints in s . The clock intervals in s indicate which clock intervals are currently active and represent witnessing intervals for type-1 and type-2 formulas; the clocks in s indicate which clocks are currently active and represent proof obligations for type-3 and type-4 formulas.

The transitions of \mathcal{M}_ϕ are all triples $s \xrightarrow{\lambda} s'$ that satisfy certain global consistency criteria. Both the local and the global consistency conditions are defined in the following catalog. For every state $s \subseteq Closure(\phi)$ and every transition $s \xrightarrow{\lambda} s'$ with source state s :

Logical consistency

- For each atomic proposition $p \in P$, precisely one of p and $\neg p$ is in s .
- If the formula $\psi_1 \wedge \psi_2$ is in s , then both ψ_1 and ψ_2 are in s .

- If the formula $\psi_1 \vee \psi_2$ is in s , then either ψ_1 or ψ_2 is in s .

These conditions ensure that no state contains subformulas of ϕ that are mutually inconsistent.

Timing consistency

- s contains at most one of the clock constraints $0 < (K - C)$, $0 \subset (K - C)$, $0 = (K - C)$, and $(K - C) = \emptyset$ for each clock interval C . Furthermore, no two clock intervals in s share clocks; for instance, s does not contain both the clock intervals (x, y) and $[x, y)$.
- s contains at most one of the clock constraints $x_\psi \in I$ and $x_\psi > I$ for each type-3 or type-4 formula ψ .
- If s contains $x_{sing} = 0$, then $x_{sing} \notin \lambda$. If s does not contain $x_{sing} = 0$, then $x_{sing} \in \lambda$ and s' contains $x_{sing} = 0$.

These conditions guarantee that no state contains clock constraints that are mutually inconsistent. We say that a state s is *singular* iff it contains $x_{sing} = 0$; otherwise s is *open*. The last clause of the above conditions ensures that singular and open states alternate along any run.

Type- 1 formulas

Consider a type-1 formula $\psi = \psi_1 \mathcal{U}_I \psi_2$ in the closure set.

Firstly, if ψ is in s , then there is some clock interval $C = C_j(\psi)$ such that

- $(K - C) \cap I \neq \emptyset$ is in s , and
- either C is in s , or s is singular and C is in s' and the clocks associated with C are not in λ .

The first condition checks that the interval $K-C$ is an appropriate candidate for witnessing the formula ψ . The second condition activates the clock interval C to represent a witnessing interval for ψ .

Secondly, if some clock interval $C = C_j(\psi)$ is in s , then

- if either $0 = (K - C)$ or $0 \subset (K - C)$ is in s , then ψ_2 is in s , and
- if either $0 < (K - C)$ or $0 \subset (K - C)$ is in s , then ψ_1 is in s , and

- the clocks associated with C are not in λ and either C or $(K - C) = \emptyset$ is in s' .

The first two conditions verify that the active clock interval C represents indeed a witness for the formula ψ . The final condition keeps the clock interval C active as long as necessary.

Suppose that these conditions are satisfied along a run r and the formula ψ is in a state at time t . Also assume (the induction hypothesis) that, along the run r , whenever a state at time t' contains a subformula ψ' of ψ , then $\rho_r^{t'} \models \psi'$. A clock interval $C = C_j(\psi)$ is activated at time t . It is not hard to show that the interval $t + K - C$ is a witnessing interval for ψ under ρ_r^t . By Lemma 4.1, it follows that $\rho_r^t \models \psi$.

Conversely, if $\rho^t \models \psi$, then there is a run r that satisfies all conditions. This is because, by Lemma 4.2, the automaton can, at time t , either share an already activated clock interval $C_j(\psi)$ or has enough clocks to activate an unused clock interval $C_j(\psi)$.

Type-2 formulas

Consider a type-2 formula $\psi = \psi_1 \text{ } I \cup \psi_2$ in the closure set.

Firstly, if ψ is in s , then either

- $\diamond_{(0,\infty) \cap (<I)} \psi_2$ is in s

or there is some clock interval $C = C_j(\psi)$ such that

- $I \subseteq (K - C)$ is in s , and
- either C is in s , or s is singular and C is in s' and the clocks associated with C are not in λ .

If $\diamond_{(0,\infty) \cap (<I)} \psi_2$ holds then so does ψ . The second clause corresponds to guessing a witness. The first condition checks that the interval $K - C$ is an appropriate candidate for witnessing the formula ψ . The second condition activates the clock interval C to represent a witnessing interval for ψ .

Secondly, if some clock interval $C = C_j(\psi)$ is in s , then

- if either $0 = (K - C)$ or $0 \subset (K - C)$ is in s , then ψ_1 is in s , and
- either ψ_2 is in s , or the clocks associated with C are not in λ and either C or $(K - C) = \emptyset$ is in s' .

These conditions ensure that the active clock interval C represents indeed a witness for the formula ψ and that it is kept active as long as necessary.

Soundness and completeness of these conditions follow by the Lemmas 4.1 and 4.3.

Type-3 formulas

Consider a type-3 formula $\psi = \diamond_I \psi'$ in the closure set

Firstly, if ψ is in s , then either

- s is singular and $\mathbf{x}_\psi \in s'$, or
- s is open and I is right-open and $\hat{\psi}$ is in s , or
- s is open and I is right-closed and \mathbf{x}_ψ is in s .

These conditions activate a clock to represent a proof obligation. Lemma 4.4 justifies the decision to check, if s is open, instead of ψ the weaker type-3 formula $\hat{\psi}$.

Secondly, if \mathbf{x}_ψ is in s , then

- $\mathbf{x}_\psi \in I$ is in s , and
- either ψ' is in s , or \mathbf{x}_ψ is in s' and $\mathbf{x}_\psi \notin \lambda$.

These conditions verify the proof obligation that is represented by the clock \mathbf{x}_ψ and keep it active as long as necessary.

Type-4 formulas

Consider a type-4 formula $\psi = \square I s'$ in the closure set.

Firstly, if ψ is in s , then either

- s is singular and $\mathbf{x}_\psi \in s'$ and $\mathbf{x}_\psi \in \lambda$, or
- s is open and I is right-closed and $\hat{\psi}$ is in s , or
- s is open and I is right-open and $\mathbf{x}_\psi \in s$ and $\mathbf{x}_\psi \in s'$ and $\mathbf{x}_\psi \in \lambda$.

These conditions activate a clock to represent a proof obligation, and reset it, as was justified in the previous subsection. Recall that if s is open, then instead of checking ψ , it suffices to check the weaker type-4 formula $\hat{\psi}$.

Secondly, if \mathbf{x}_ψ is in s then

- ψ' is in s , and
- either x_ψ or $x > I$ is in s' .

The first condition verifies the proof obligation that is represented by the clock x_ψ , and the second condition keeps it active as long as necessary.

Type-5 formulas

Consider a type-5 formula $\psi = \psi_1 \mathcal{U} \psi_2$ in the closure set. Whenever ψ is in s , then either

- s is singular and $\psi \in s'$, or
- s is open and ψ_1 is in s , and either ψ_2 is in s or ψ_2 is in s' or both ψ_1 and ψ are in s' .

These conditions ensure that unconstrained *until* formulas are propagated correctly (remember that singular and open intervals alternate).

Type-6 formulas

Consider a type-6 formula $\psi = \bullet I \psi'$ in the closure set. Whenever ψ is in s , then either

- s is singular and $\psi \in s'$, or
- s is open and $\psi' \in s$ and both ψ' and ψ are in s' .

These conditions guarantee that unconstrained *always* formulas are propagated forever.

This concludes the definition of the timed automaton \mathcal{M}_ϕ . The runs of \mathcal{M}_ϕ are defined as before. We put, however, additional fairness requirements on the timed state sequences that are generated by \mathcal{M}_ϕ . A run r is called *accepting* iff for every type-5 formula ψ of the form $\psi_1 \mathcal{U} \psi_2$, if ψ is in some state s along r , then ψ_2 is in some later state s' .

The following main lemma states the correctness of our construction by relating the accepting runs of \mathcal{M}_ϕ to the models of ϕ .

Lemma 4.5 (Correctness of \mathcal{M}_ϕ) *A timed state sequence ρ satisfies an MITL-formula ϕ iff the timed automaton \mathcal{M}_ϕ has an accepting run r with $\rho = \rho_r$.*

Proof of Lemma 4.5 It can be shown, by induction on the structure of ϕ , that given an accepting run r of \mathcal{M}_ϕ , if a subformula ψ of ϕ is in a state s in r at time $t \in \mathbb{R}^+$, then $\rho_r^t \models \psi$. We have outlined the crucial arguments for the six interesting cases of temporal subformulas above.

Conversely, given a δ -fine model ρ of ϕ with alternating singular and open intervals, we can construct an accepting run r of \mathcal{M}_ϕ such that $\rho = \rho_r$. The Lemmas 4.2 and 4.3 instruct us how to use the limited number of available clocks to mark witnessing intervals. ■

This result yields algorithms for checking the satisfiability and validity of the given MITL-formula ϕ . To check satisfiability, we first construct the timed automaton \mathcal{M}_ϕ , and then we use the algorithm that checks whether \mathcal{M}_ϕ has any accepting run to test if ϕ has a model. Similarly, ϕ is valid iff $\mathcal{M}_{\neg\phi}$ has no accepting run.

4.7 Complexity of MITL

We conclude this section by showing that our decision procedure for MITL is in EXPSPACE, and that this is optimal, because the decision problem for MITL is EXPSPACE-complete.

Recall that the size $|Closure(\phi)|$ of the closure set of ϕ is $O(N \cdot K)$, where N is the number of atomic propositions, boolean connectives, and temporal operators in ϕ , and $K - 1$ is the product of the largest constant in ϕ and the least common denominator of all constants in ϕ . Clearly, $|Closure(\neg\phi)| = O(N \cdot K)$ as well.

Hence the number of states in \mathcal{M}_ϕ and $\mathcal{M}_{\neg\phi}$ is $O(2^{N \cdot K})$. Consequently, the description of $\mathcal{M}_{\neg\phi}$ can be given in space polynomial in $N \cdot K$; that is, in space exponential in the length of ϕ , assuming binary encoding of all interval end-points. The emptiness problem for a timed automaton \mathcal{M} can be solved in space polynomial in the length of the description of \mathcal{M} . It follows that the validity of ϕ can be decided in space polynomial in $N \cdot K$, that is, in EXPSPACE.

The lower bound of EXPSPACE for MITL can be shown along the lines of the proof of the EXPSPACE-hardness of the real-time logic MTL [AH90].

Theorem 4.1 (Complexity of MITL) *The decision problem of MITL is EXPSPACE-complete. Furthermore, we have an EXPSPACE algorithm that solves this problem.*

5 Model Checking

Model checking is a powerful and well-established technique for the automatic verification of finite-state systems (see, for example, [BCM⁺90]); it compares a temporal-logic specification of a system against a state-transition description of the system.

In the qualitative case, the system is modeled by its state-transition graph, also known as Kripke structure, and the specification may be presented as a formula of the propositional linear temporal logic PTL [LP84]. For real-time systems, model checking algorithms have been developed for linear temporal logics under a digital-clock interpretation of time [AH89, AH90, HLP90] as well as for branching-time logics under a continuous interpretation of time [ACD90, Lew90]. Using our results about MITL, we can present a real-time verification procedure that checks *linear* specifications under a *continuous* model of time.

We model a real-time system by a timed automaton M and give the specification as a formula ϕ of MITL. Hence the *model checking* problem is to decide whether or not the automaton M satisfies the specification ϕ :

$$\mathcal{M} \stackrel{?}{\models} \phi$$

Our construction for testing the satisfiability of MITL-formulas can be used to develop an algorithm for model checking. The first step is to construct a timed automaton $\mathcal{M}_{\neg\phi}$ such that its accepting runs precisely capture the models of the negated formula $\neg\phi$: for every timed state sequence ρ , $\mathcal{M}_{\neg\phi}$ has an accepting run τ with $\rho_\tau = \rho$ iff $\rho \models \neg\phi$.

The model checking question can, then, be reformulated as follows: $M \models \phi$ iff no timed state sequence is generated by both M and $\mathcal{M}_{\neg\phi}$. The next step in the model checking algorithm is to construct a timed automaton M' that is the product of M and $\mathcal{M}_{\neg\phi}$; a timed state sequence is generated by M' iff it is generated by both M and $\mathcal{M}_{\neg\phi}$.

The *product* construction for timed automata presented in [AD90] can be easily modified to our version of timed automata. We assume that the clock sets of the component automata, M and $M_{\neg\phi}$, are disjoint. The set of clocks of M' is the union of the clocks of the component automata. The states of M' are of the form $\langle s, s' \rangle$, where s is a state of M and s' is a state of $\mathcal{M}_{\neg\phi}$ and both s and s' agree on the assignment of truth values to propositions. The clock constraints for $\langle s, s' \rangle$ are the conjunctions of the clock constraints for s and s' . For any pair of transitions $u \xrightarrow{\lambda} v$ and $u' \xrightarrow{\lambda'} v'$

in M and $\mathcal{M}_{\neg\phi}$, respectively, the product automaton has *three* transitions: $\langle u, u' \rangle \xrightarrow{\lambda \cup \lambda'} \langle v, v' \rangle$, $\langle u, u' \rangle \xrightarrow{\lambda} \langle v, u' \rangle$, and $\langle u, u' \rangle \xrightarrow{\lambda'} \langle u, v' \rangle$. Thus the transitions of M' simulate the joint behavior of the two component automata. The acceptance conditions of the individual automata are handled as in the product construction for (untimed) w-automata.

Hence we have reduced the model checking problem to the emptiness question for timed automata: $M \models \phi$ iff M' has no accepting runs. The size of M' is polynomial in the sizes of M and $\mathcal{M}_{\neg\phi}$. Consequently, the description of M' is exponential in the length of ϕ , and polynomial in the length of the description of M . Since the emptiness for timed automata can be solved in PSPACE, it follows that the model checking problem can be solved in EXPSPACE.

As for all linear temporal logics, the model checking question for MITL is no simpler than the satisfiability question: a formula ϕ is unsatisfiable iff the universal timed automaton, which generates all possible timed state sequences, satisfies $\neg\phi$. Thus EXPSPACE-hardness of satisfiability implies EXPSPACE-hardness of model checking. The following theorem follows:

Theorem 5.1 (Model checking) *The problem of checking if a timed automaton M satisfies an MITL-formula ϕ is EXPSPACE-complete.*

The time complexity of the model checking algorithm is polynomial in the qualitative part of the system description, exponential in the qualitative part of the MITL-specification, exponential in the timing part of the system description, and doubly exponential in the timing part of the specification. Compared to this the model checking algorithm for PTL [LP84] is polynomial in the size of the Kripke structure and exponential in the size of the specification.

Thus moving to *real-time* gives an additional *exponential* blow-up. This blow-up seems, however, unavoidable for formalisms for quantitative reasoning about time. It occurs even in the simplest case — synchronous processes that are clocked by a digital clock — in which we can model time by a discrete domain and identify next-state with next-time [EMSS89, AH90].

Acknowledgements. We thank Moshe Vardi for encouraging us repeatedly to look for logics of continuous time, Ron Koymans for helpful suggestions, and David Dill and Zohar Manna for their guidance.

References

- [ACD90] Rajeev Alur, Costas Courcoubetis, and David L. Dill. Model-checking for real-time systems. In *Proceedings of the Fifth Annual IEEE Symposium on Logic in Computer Science, 1990*.
- [AD90] Rajeev Alur and David L. Dill. Automata for modeling real-time systems. In *17th International Colloquium on Automata, Languages, and Programming*. Springer-Verlag Lecture Notes in Computer Science 443, 1990.
- [AH89] Rajeev Alur and Thomas A. Henzinger. A really temporal logic. In *Proceedings of the 30th Annual IEEE Symposium on Foundations of Computer Science, 1989*.
- [AH90] Rajeev Alur and Thomas A. Henzinger. Real-time logics: complexity and expressiveness. In *Proceedings of the Fifth Annual IEEE Symposium on Logic in Computer Science, 1990*.
- [BCM⁺90] J. R. Burch, E. M. Clarke, K. L. McMillan, D. L. Dill, and L. J. Hwang. Symbolic model checking: 10^{20} states and beyond. In *Proceedings of the Fifth Annual IEEE Symposium on Logic in Computer Science, 1990*.
- [EMSS89] E. Allen Emerson, Aloysius K. Mok, A. Prasad Sistla, and Jai Srinivasan. Quantitative temporal reasoning. Presented at the First Annual Workshop on Computer-aided Verification, Grenoble, France, 1989.
- [HLP90] Eyal Harel, Orna Lichtenstein, and Amir Pnueli. Explicit-clock temporal logic. In *Proceedings of the Fifth Annual IEEE Symposium on Logic in Computer Science, 1990*.
- [JM86] Farnam Jahanian and Aloysius K. Mok. Safety analysis of timing properties in real-time systems. *IEEE Transactions on Software Engineering*, SE-12, 1986.
- [Koy90] Ron Koymans. Specifying real-time properties with metric temporal logic. *Journal of Real-time Systems*, 2, 1990.
- [Lew90] Harry R. Lewis. A logic of concrete time intervals. In *Proceedings of the Fifth Annual IEEE Symposium on Logic in Computer Science, 1990*.

- [LP84] Orna Lichtenstein and Amir Pnueli. Checking that finite-state concurrent programs satisfy their linear specification. In *Proceedings of the 11 th Annual ACM Symposium on Principles of Programming Languages, 1984*.
- [Ost90] Jonathan S. Ostroff. *Temporal Logic of Real-time Systems*. Research Studies Press, 1990.
- [Rog67] Hartley Rogers, Jr. *Theory of Recursive Functions and Effective Computability*. McGraw-Hill, 1967.
- [Tho90] Wolfgang Thomas. Automata on infinite objects. In Jan van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B. Elsevier, 1990.