Metric Linear-Time Temporal Logic with Strict First-Time Semantics

Eric Alsmann
Theoretical Computer Science / Formal Methods, University of Kassel, Germany

Martin Lange
Theoretical Computer Science / Formal Methods, University of Kassel, Germany

Abstract

We introduce strict first-time semantics for the Until operator from linear-time temporal logic which makes assertions not just about some future moment but about the first time in the future that its argument should hold. We investigate Metric Linear-Time Temporal Logic under this interpretation in terms of expressive power, relative succinctness and computational complexity. While the expressiveness does not exceed that of pure LTL, there are properties definable in this logic which can only be expressed in LTL with exponentially larger formulas. Yet, we show that the complexity of the satisfiability problem remains PSPACE-complete which is in contrast to the EXPSPACE-completeness of Metric LTL. The motivation for this logic originates in a study of the expressive power of State Space Models, a recently proposed alternative to the popular transformer architectures in machine learning.

2012 ACM Subject Classification Theory of computation → Modal and temporal logics; Theory of computation → Automata over infinite objects

Keywords and phrases linear-time temporal logic, metric temporal logic, computational complexity, Streett automata

Digital Object Identifier 10.4230/LIPIcs.TIME.2025.3

1 Introduction

The linear-time temporal logic LTL is a well-known formalism for specifying properties of runs of reactive systems [14]. Its model-theoretic and computational properties are well researched and understood, for instance its satisfiability problem being PSPACE-complete [17] and its expressiveness coinciding with that of First-Order Logic [6], resp. star-free regular expressions over ω-words or over finite words [7].

The relative weakness in expressive power has led to the study of several extensions of LTL, some of which genuinely extend its expressiveness, typically to that of full (ω-)regularity, for instance the Linear-Time μ-Calculus [3, 19], the industry standard PSL [5], Quantified LTL [16], etc. Others extend LTL only pragmatically by providing further constructs without extending the expressive power altogether, but providing means to express certain properties more easily.

One such variant is the Metric Linear-Time Temporal Logic MTL [1]. The term "metric temporal logic" does not uniquely identify one particular formalism, not even within the linear-time framework. It describes temporal logics in which the temporal operators are extended so that they do not simply make assertions about the future (or past) moments in a linear sequence of events, but additionally constrain their distance to the current moment. For example $p\,\mathtt{U}_{[3,5]}\,q$ does not only demand that q holds at some point in the future with p being continuously true up to that point, but this future point also has to occur within three to five time units from now on. The exact semantics then depends on the underlying models etc. Metric temporal logic has been developed for reasoning about real-time systems [9], but

it can also be used for systems evolving in discrete steps. Here we use MTL only as a logic over discrete time, resp. (untimed) ω-words. Consequently, the intervals adorning the metric operators are always interpreted over the natural numbers.

It is not hard to see that MTL does not exceed LTL in expressiveness: temporal operators with metric constraints can easily be unravelled. The property above can be expressed as "*p holds now and in the next two steps; afterwards we have q or p followed by q or p followed by another p, followed by q*" using only the next-time operator (in this case). A general translation from MTL to LTL is an easy exercise, and it incurs a blow-up that is polynomial in the value of the involved interval bounds. Hence, when such bounds are encoded unarily, MTL satisfiability is also PSpace-complete. The more reasonable assumption of binary encodings then yields an exponential translation into LTL and therefore an ExpSpace upper bound on its satisfiability problem. This is tight, satisfiability for MTL with binarily encoded interval bounds is in fact ExpSpace-complete [10].

The lower bound uses a standard reduction from the word problem for exponentially space bounded Turing Machines (TM). The key to the proof is then MTL's ability to express "something holds after 2^n steps" which, using binary encodings, can be done with a formula of size polynomial in n. This is used to express, for instance, that a symbol b is seen at that distance, i.e. at the same cell on the tape in the next configuration, formalising a potential writing operation of the TM. Clearly, the tape may contain more occurrences of that symbol, so the one seen at distance 2^n is usually not the first one to be seen in the future.

This paper is motivated by the study of the expressive power of a recent machine learning model, so-called State-Space Models (SSM) [13]. They have been proposed as an alternative to the well-known transformer architectures underlying prominent Large Language Model tools, promising more efficient evaluation of long input strings. Here we do not go further into the details of SSMs; they merely serve as a motivation for studying a particular variant of MTL which is geared towards the expressive capabilities of SSMs. To motivate this variant, we recall the well-known unfolding principles for temporal operators.

▶ **Example 1.** Suppose that, in the context of some decision procedure, we were to establish the truth of the LTL formula $\mathtt{G}(p\,\mathtt{U}\,q)$ at position i of some ω-sequence of events. By unfolding the $\mathtt{G}$-operator, this typically amounts to establishing truth of both (i) $p\,\mathtt{U}\,q$ and (ii) $\mathtt{X}\,\mathtt{G}(p\,\mathtt{U}\,q)$ at moment i. To solve (i) we can either prove q or defer it to a later point by unfolding the Until-formula and proving p as well as (iii) $\mathtt{X}(p\,\mathtt{U}\,q)$ at moment i. Now, (ii) and (iii) both start with a Next-operator, so they imply proving $\mathtt{G}(p\,\mathtt{U}\,q)$ and (iii') $p\,\mathtt{U}\,q$ at moment $i+1$. The former can be handled as done at moment i; in particular it means proving $p\,\mathtt{U}\,q$ at moment $i+1$ which is already the task identified as (iii'). In a sense, there are two reasons for the need to prove $p\,\mathtt{U}\,q$ at moment $i+1$, and they just collapse to a single task.

Now suppose that the original formula was one of MTL, namely $\mathtt{G}(p\,\mathtt{U}_{[3,5]}\,q)$. Here the reasoning can be done in the same way, but when unfolding the metric Until we obtain an index shift: $p\,\mathtt{U}_{[3,5]}\,q$ is established at moment i when $\mathtt{X}(p\,\mathtt{U}_{[2,4]}\,q)$ is established there. Hence, when carrying out this kind of reasoning we obtain two tasks regarding moment $i+1$ that do not collapse into one, namely to prove $p\,\mathtt{U}_{[2,4]}\,q$ as well as $p\,\mathtt{U}_{[3,5]}\,q$ there.

This can also be seen as an indication for MTL's higher complexity: truth of a formula at moment i does not just depend on the truth of its subformulas at moments i and $i+1$, but also on variants of the subformulas obtained through index shifts.

The mechanisms in SSMs naturally suggest a comparison of their expressiveness with formulas of *metric* linear-time temporal logic (on finite words), since SSMs can easily track the distance between occurrences of events in a word, they seem to be unable to cope with the blow-up in formulas that need to be tracked through unfolding.

In this paper we initiate the study of a variant of MTL– called *Metric Temporal Logic with Strict First-Time Semantics* here, MTL^1– which is defined because of a certain similarity to the mechanics in SSMs. It features the Metric Until operator $\varphi \,\mathtt{U}^1_{[x,y]}\, \psi$ which does not just demand that some future moment at distance between x and y satisfies ψ, but this also has to be the *first time* (from the current moment on) that ψ is satisfied. Note that this remedies the problem with a potential blow-up indicated by the example above: while $(p \,\mathtt{U}_{[2,4]}\, q) \wedge (p \,\mathtt{U}_{[3,5]}\, q)$ cannot be simplified into a formula using only a single metric Until operator, $(p \,\mathtt{U}^1_{[2,4]}\, q) \wedge (p \,\mathtt{U}^1_{[3,5]}\, q)$ is indeed equivalent to $p \,\mathtt{U}^1_{[3,4]}\, q$, suggesting that the strict first-time semantics makes formulas easier to handle computationally. It is also not hard to see that the proof of ExpSpace-hardness for MTL breaks down under the strict first-time semantics, as suggested above. So apart from the obvious question about MTL^1's expressiveness in relationship to the other variants of LTL, its complexity somewhere between PSpace and ExpSpace is to be pinpointed as well. An ExpSpace upper bound is obtained easily because of a simple linear translation into MTL, based on principles like $\varphi \,\mathtt{U}^1_{[x,y]}\, \psi \equiv (\varphi \wedge \neg\psi) \,\mathtt{U}_{[x,y]}\, \psi$.

This paper is organised as follows. In Section 2 we formally recall known concepts needed for this study, in particular LTL and its extension MTL. We also recall the definition of Streett automata [18] which will serve as a main tool in a decision procedure for MTL^1 which is formally introduced in Section 3. There, we also investigate its expressiveness and succinctness relative to LTL. In Section 4 we construct a singly exponential translation from MTL^1 into Streett automata thus pinpointing the former's complexity as being PSpace-complete only. In Section 5 we conclude with remarks on further work in this area.

2 Preliminaries

We recall the necessary definitions from temporal logics and automata theory.

Linear-time temporal logic. Let $\mathcal{P} = \{p, q, \ldots\}$ be a non-empty, countable set of atomic propositions. Formulas of the linear-time temporal logic LTL are given by the following grammar.

$$\varphi \;::=\; q \mid \varphi \wedge \varphi \mid \neg\varphi \mid \mathtt{X}\,\varphi \mid \varphi \,\mathtt{U}\, \varphi$$

where $q \in \mathcal{P}$. Further Boolean operators like $\mathtt{tt}, \mathtt{ff}, \vee, \rightarrow, \leftrightarrow$ are introduced as abbreviations in the usual way, and so are the temporal operators $\mathtt{F}\,\varphi := \mathtt{tt}\,\mathtt{U}\,\varphi$, $\mathtt{G}\,\varphi := \neg\,\mathtt{F}\,\neg\varphi$, $\varphi\,\mathtt{R}\,\psi := \neg(\neg\varphi\,\mathtt{U}\,\neg\psi)$.

The set $\mathit{Sub}(\varphi)$ of subformulas of φ is defined in the usual way. The size of a formula φ is measured in terms of the number of its subformulas: $|\varphi| := |\mathit{Sub}(\varphi)|$. For our purposes here, it makes no difference whether formulas with abbreviated operators are written out in the original syntax or the abbreviated operators are seen as first-class citizens. The size in the former case is only larger by a constant factor compared to the latter case.

Formulas of LTL are interpreted over traces of labelled transition systems which are just ω-words over the alphabet $2^{\mathcal{P}}$. Note that any formula can contain only finitely many propositional symbols. Hence, the alphabet underlying any given formula can always be assumed to be finite. Let $w = a_0 a_1 \ldots \in (2^{\mathcal{P}})^\omega$ and $i \in \mathbb{N}$. Satisfaction of an LTL formula φ in w at position i is explained inductively as follows.

$$\begin{aligned} w, i \models q \quad & \text{iff} \quad q \in a_i \\ w, i \models \varphi \wedge \psi \quad & \text{iff} \quad w, i \models \varphi \text{ and } w, i \models \psi \end{aligned}$$

$$
\begin{array}{lll}
w,i \models \neg\varphi & \text{iff} & w,i \not\models \varphi \\
w,i \models \mathtt{X}\,\varphi & \text{iff} & w,i+1 \models \varphi \\
w,i \models \varphi\,\mathtt{U}\,\psi & \text{iff} & \text{there is } j \geq i \text{ s.t. } w,j \models \psi \text{ and } w,h \models \varphi \text{ for all } h \text{ with } i \leq h < j
\end{array}
$$

A word w satisfies a formula φ, written $w \models \varphi$, if φ is satisfied in its initial position, i.e. $w,0 \models \varphi$. The *language* of a formula φ is the set of all words satisfying it: $L(\varphi) := \{w \mid w \models \varphi\}$. Two formulas are equivalent, written $\varphi \equiv \psi$, when their languages coincide, i.e. $L(\varphi) = L(\psi)$. Note that this is the case if and only if they are satisfied by exactly the same positions in the same words, i.e. $w,i \models \varphi$ iff $w,i \models \psi$. The "if"-direction is trivial because two formulas that are satisfied by the same words at arbitrary positions are clearly satisfied by the same words at initial positions. The "only if"-direction holds because LTL has future-only operators, so we have $a_0a_1\ldots,i \models \varphi$ iff $a_ia_{i+1}\ldots,0 \models \chi$ for any φ. Thus, any position in a word is the initial position of another word, namely the suffix starting at this position.

A formula is *valid*, written $\models \varphi$, if $L(\varphi) = (2^{\mathcal{P}})^\omega$, i.e. it is satisfied by any word. The satisfiability problem – given a formula φ, decide whether $L(\varphi) \neq \emptyset$ – is well-known to be decidable. We write $\mathrm{SAT}(\mathcal{L})$ for the satisfiability problem for logic $\mathcal{L}$. We assume familiarity with standard complexity classes, in particular the space complexity classes NLogSpace, PSpace and ExpSpace. For further details we refer to standard textbooks in complexity theory, e.g. [2].

▶ **Proposition 2** ([17]). *SAT(LTL) is PSpace-complete.*

Metric linear-time temporal logic. The term "metric" temporal logic usually refers to extensions of ordinary temporal logics (like LTL) with modifications of the temporal operators that impose restrictions on the moments that witness their satisfaction in case of an Until or a Finally, resp. relax the future moments under consideration of a Generally, resp. Release formula. The logic MTL is obtained by extending the syntax of LTL with a metric version of Until, written $\varphi\,\mathtt{U}_I\,\psi$ for a non-empty interval I over the natural numbers.

We use standard notation for closed, half-open and open intervals like $[x,y], (x,y], [x,y)$ and (x,y) for $x,y \in \mathbb{N} \cup \{\infty\}$ with $x \leq y$. Because of the discrete nature of $\mathbb{N}$ and the lack of a direct predecessor of ∞ in this structure, we can restrict our attention to intervals of the form $[x,y]$ for $x,y \in \mathbb{N}$, and of the form $[x,\infty)$ for $x \in \mathbb{N}$. Note that $(x,y] = [x+1,y]$ etc.

Intervals of the form $[0,y]$ or $[x,\infty)$ for some $x,y \in \mathbb{N}$ are called *simple* and written more compactly as ${\leq}y$, resp. ${\geq}x$.

The semantics of formulas of MTL is extended accordingly for words $w = a_0a_1\ldots \in (2^{\mathcal{P}})^\omega$ and $i \in \mathbb{N}$ by

$$
\begin{array}{lll}
w,i \models \varphi\,\mathtt{U}_I\,\psi & \text{iff} & \text{there is } j \geq i \text{ s.t. } j-i \in I \text{ and } w,j \models \psi \text{ and} \\
& & w,h \models \varphi \text{ for all } h \text{ with } i \leq h < j
\end{array}
$$

Hence, $\varphi\,\mathtt{U}_I\,\psi$ note only requires a position i in w to satisfy $\varphi\,\mathtt{U}\,\psi$ in the usual sense that some future moment satisfies ψ and all future moments before that satisfy φ. It additionally requires that future moment to be found at a distance which falls into the interval I.

This clearly extends the non-metric version of LTL because $\varphi\,\mathtt{U}\,\psi \equiv \varphi\,\mathtt{U}_{\geq 0}\,\psi$. The metric extension is then applied to the other derived temporal operators $\mathtt{F}, \mathtt{G}$ and $\mathtt{R}$ accordingly.

Because of the use of natural numbers as interval bounds in temporal operators, number of subformulas is not a realistic measure of the representation size of a formula anymore. Note that $\mathtt{F}_{[x,y]}\,p$, stating that p occurs in a word at some distance between x and y would

have constant size regardless of the values of x and y. We assume that such interval bounds are given in binary encoding and therefore measure the size of a formula more appropriately as follows. Let $m(\varphi) :=$

$$\max\{k \in \mathbb{N} \mid \exists \psi_1, \psi_2 \in \mathit{Sub}(\varphi), \exists h \in \mathbb{N} \text{ s.t. } \psi_1 \,\mathtt{U}_I\, \psi_2 \in \mathit{Sub}(\varphi) \text{ for } I = [h,k] \text{ or } I = [k,\infty)\}$$

denote the largest integer constant that is used as an interval bound in a subformula of φ. Then $|\varphi| := |\mathit{Sub}(\varphi)| \cdot \lceil \log m(\varphi) \rceil$.

▶ **Proposition 3** ([10]). *SAT(MTL) is ExpSpace-complete.*

Finite automata on ω-words. Let $\mathcal{P}$ be a finite set of atomic propositions. A nondeterministic Streett automaton (NSA) is an $\mathcal{A} = (Q, \mathcal{P}, q_I, \delta, \mathcal{F})$ where Q is a finite set of states, $q_I \in Q$ is the designated initial state, $\delta \subseteq Q \times \mathbb{B}(\mathcal{P}) \times Q$ is a finite set of transition triples (p, β, q) with $\mathbb{B}(\mathcal{P})$ denoting the set of Boolean formulas over $\mathcal{P}$ (with the usual Boolean operators $\wedge$, $\vee$, $\neg$, $\ldots$). Finally, $\mathcal{F} = \{(F_1, G_1), \ldots, (F_m, G_m)\}$ with $F_i, G_i \subseteq Q$ for all $i = 1, \ldots, m$ is the acceptance condition.

Note that here such finite automata have a symbolically represented transition table instead of the general $\delta \subseteq Q \times \Sigma \times Q$ for a finite alphabet Σ. The reason for this choice is their use in a decision procedure for a temporal logic whose formulas are naturally interpreted over ω-words whose letters are finite sets of propositions, i.e. $\Sigma = 2^{\mathcal{P}}$. Instead of enumerating all such possible sets and explaining the automaton's one-step behaviour for each step, the symbolic representation here is more convenient for such special cases. There are straightforward translations for the transition relations between symbolic and explicit representations, and for a fixed set of propositions, they are polynomial. Hence, we can safely use Streett automata in this form and still appeal to known results about them, even though they were perhaps formulated for the form with an explicit alphabet.

A run of the NSA $\mathcal{A} = (Q, \mathcal{P}, q_I, \delta, \mathcal{F})$ with $\mathcal{F} = \{(F_1, G_1), \ldots, (F_m, G_m)\}$ on a word $w = a_0 a_1 \ldots \in (2^{\mathcal{P}})^\omega$ is a $\rho = q_0, q_1, \ldots \in Q^\omega$ such that $q_0 = q_I$ and for all $i \in \mathbb{N}$ there is some $(p, \beta, q) \in \delta$ such that $q_i = p$, $a_i \models \beta$ and $q = q_{i+1}$. Here, satisfaction of a Boolean formula β over $\mathcal{P}$ by a set $a \subseteq \mathcal{P}$, $a \models \beta$, is explained in the usual way: β evaluates to true when all $q \in a$ are set to true and all $q \notin a$ are set to false.

We write $\mathit{Inf}(\rho)$ to denote the set of all states that occur infinitely often in ρ. By finiteness of Q, we always have $\mathit{Inf}(\rho) \neq \emptyset$.

The run ρ is accepting if it satisfies the Streett condition $\mathcal{F}$ in the sense that for all $i = 1, \ldots, m$: if $\mathit{Inf}(\rho) \cap F_i \neq \emptyset$ then $\mathit{Inf}(\rho) \cap G_i \neq \emptyset$. As usual, $L(\mathcal{A})$ is the language of the NSA $\mathcal{A}$, and it consists of all words on which $\mathcal{A}$ has an accepting run.

So both NSA and formulas of linear-time temporal logics defines languages of ω-words which is why these different formalisms can be compared to one another in terms of expressiveness, and satisfiability of a formula corresponds to non-emptiness of the language of an automaton. Algorithms for non-emptiness problems for ω-automata are routinely used to obtain decision procedures for linear-time temporal logics [21, 20, 4]. It has been shown that non-emptiness for Streett automata can be decided in polynomial time [8, 12] and this requires an explicit construction. A polynomial (equivalence-preserving) translation into Büchi automata is not possible, let alone one computable in logarithmic space. This would immediately transfer the upper bound of NLogSpace to Streett automata. Nevertheless, it does hold, too; it simply needs to be shown directly.

▶ **Theorem 4.** *The non-emptiness problem for NSA is decidable in NLogSpace.*

Proof. The key observation is the following. The language of an NSA $\mathcal{A} = (Q, \mathcal{P}, q_I, \delta, \mathcal{F})$ with $n := |Q|$ and $\mathcal{F} = \{(F_1, G_1), \ldots, (F_k, G_k)\}$ is non-empty iff there is an ultimately periodic word $w = uv^\omega \in L(\mathcal{A})$. This follows from finiteness of n and k by the pigeon hole principle. An accepting run on an arbitrary word must eventually traverse a state q for the second time such that in between, for every $i = 1, \ldots, k$, either no state from F_i or some state from G_i has been seen.

This gives rise to a nondeterministic algorithm for deciding non-emptiness. It guesses, step-by-step, the states of a run and also nondeterministically remembers some state q occurring in this simulation. It then maintains $2k$ bits to remember, for each $i = 1, \ldots, k$, whether some state in F_i and some state in G_i has been seen. It accepts, when q occurs again, and the bits indicate that the Streett condition has been met in between.

In order to terminate and reject on computation paths with unsuccessful guesses, it counts the number of steps done in this simulation. It is not hard to see that the first occurrence of q can be required to occur after at most n steps. The second occurrence can be expected to occur after no more than a further $2nk$ steps, for otherwise the run contains parts that could be skipped. Hence, the space needed for the counter is at most logarithmic in $|\mathcal{A}|$. ◀

3 Metric LTL with First-Time Semantics

Syntax and Semantics. We introduce *Metric Linear-Time Temporal Logic with Strict First-Time Semantics* (MTL^1) which, instead of metric Until formulas of the form $\varphi \,\mathtt{U}_I\, \psi$, features a special modification $\varphi \,\mathtt{U}^1_I\, \psi$ that is interpreted under *strict first-time semantics*. Intuitively, it does not just demand that *some* occurrence of ψ in the future happens at a distance that falls into the interval I. Instead, it requires this to be the *first* time in the future that this happens.

▶ **Definition 5.** Let $\mathcal{P} = \{p, q, \ldots\}$ be a non-empty, countable set of atomic propositions as usual. Formulas of the linear-time temporal logic MTL^1 are built according to the following grammar.

$$\varphi \;::=\; q \mid \varphi \wedge \varphi \mid \neg\varphi \mid \mathtt{X}\,\varphi \mid \varphi \,\mathtt{U}^1_I\, \varphi$$

where $q \in \mathcal{P}$ and I is an interval over $\mathbb{N}$ as discussed above.

We also introduce its fragment sMTL^1– *Simple* MTL^1– which only features simple intervals $\leq k$ or $\geq k$ in its formulas. Other Boolean and temporal operators, in particular the strict first-time variants $\mathtt{F}^1_I$, $\mathtt{G}^1_I$ and $\mathtt{R}^1_I$ are introduced as abbreviations in the usual way. The set $Sub(\varphi)$ of subformulas of an MTL^1 formula φ is defined as usual by induction over the syntax tree of the formula.

The intuition behind the restriction to first-time occurrences is made formal as follows.

▶ **Definition 6.** Let $\mathcal{P}$ be given as above and $w = a_0 a_1 \ldots \in (2^{\mathcal{P}})^\omega$. Satisfaction of an MTL^1 formula φ at a position i of the word w is explained inductively over the structure of φ as for LTL, apart from the following case.

$$w, i \models \varphi \,\mathtt{U}^1_I\, \psi \quad \text{iff} \quad \begin{array}{l} \text{there is } j \geq i \text{ s.t. } j - i \in I, \text{ and } w, j \models \psi \text{ and} \\ \qquad w, h \models \varphi \wedge \neg\psi \text{ for all } h \text{ s.t. } i \leq h < j \end{array}$$

Expressiveness. LTL trivially embeds into sMTL^1 which trivially embeds into MTL^1. The reason for this is the fact whenever an Until-formula gets satisfied *somewhere* then there is also a first time that this happens. A more interesting question concerns the other direction, and therefore also the connection to MTL. Since LTL can trivially be embedded into MTL, we immediately obtain a translation from MTL^1 into MTL, from one from MTL^1 into LTL.

We remark that the translations introduced in the following are not polynomial for formula length, measure in terms of length of string representations, but only for formula size, measured in terms of number of subformulas, as they often require the duplication of subformulas.

The first observation about expressiveness is the expressive equivalence between MTL^1 and sMTL^1. This is perhaps a little bit surprising as this principle does not apply to MTL.

▶ **Theorem 7.** *For every $\varphi \in MTL^1$ there is a $\widehat{\varphi} \in sMTL^1$ of size $\mathcal{O}(|\varphi|)$ such that $\widehat{\varphi} \equiv \varphi$.*

Proof. We can define $\widehat{\varphi}$ inductively. The only non-trivial case is that of an Until formula. Then we have

$$\widehat{\varphi\, \mathtt{U}^1_{[x,y]}\, \psi} \;:=\; (\widehat{\varphi}\, \mathtt{U}^1_{\geq x}\, \widehat{\psi}) \wedge (\widehat{\varphi}\, \mathtt{U}^1_{\leq y}\, \widehat{\psi})$$

for $x, y \in \mathbb{N}$. Clearly, this increases the number of subformulas at most by factor 3. Correctness of this translation is straightforward by inspection of the semantics. ◀

Because of Thm. 7 we can restrict our attention to sMTL^1 formulas down below as this simplifies the technical details of various constructions slightly.

One reason for considering numerical values in formulas to be represented in binary (as opposed to unary), apart from this being natural, is the fact that the expressive power of sMTL^1– and that of MTL in fact – does not exceed that of LTL. However, translations back into LTL are polynomial in the value of interval bounds only, i.e. they are in fact exponential in the size of a formula.

▶ **Theorem 8.** *For every $sMTL^1$ formula φ there is an LTL formula $\widehat{\varphi}$ such that $\widehat{\varphi} \equiv \varphi$ and $|\widehat{\varphi}| = 2^{\mathcal{O}(|\varphi|)}$.*

Proof. We define a translation $\widehat{\cdot}\colon \mathrm{sMTL}^1 \to \mathrm{LTL}$ as follows.

$$\widehat{q} := q \qquad \widehat{\varphi \wedge \psi} := \widehat{\varphi} \wedge \widehat{\psi} \qquad \widehat{\neg\varphi} := \neg\widehat{\varphi} \qquad \widehat{\mathtt{X}\,\varphi} := \mathtt{X}\,\widehat{\varphi}$$

$$\widehat{\varphi\, \mathtt{U}^1_{\leq k}\, \psi} := \widehat{\psi} \vec{\vee} \underbrace{(\widehat{\varphi} \wedge \mathtt{X}(\widehat{\psi} \vec{\vee} \ldots \vec{\vee}\, \mathtt{X}(\widehat{\psi} \vec{\vee} (\widehat{\varphi} \wedge \mathtt{X}\,\widehat{\psi}) \ldots)))}_{k \text{ occurrences of } \mathtt{X}}$$

$$\widehat{\varphi\, \mathtt{U}^1_{\geq k}\, \psi} := \underbrace{\widehat{\varphi} \wedge \neg\widehat{\psi} \wedge \mathtt{X}(\widehat{\varphi} \wedge \neg\widehat{\psi} \wedge \mathtt{X}(\ldots \mathtt{X}(\widehat{\varphi} \wedge \neg\widehat{\psi} \wedge \mathtt{X}(\varphi\, \mathtt{U}\, \psi)) \ldots))}_{k \text{ occurrences of } \mathtt{X}}$$

The translation of an operator $\mathtt{U}^1_{\sim k}$ clearly produces a formula whose size is linear in the value k, i.e. exponential in the representation size of k. Replacing the abbreviated biased disjunctions by plain Boolean formulas only incurs a further polynomial blow-up because formula size is measured in terms of number of subformulas.

Correctness of the translation is proved by a straightforward induction on the structure of φ, showing that for all $w \in (2^{\mathcal{P}})^\omega$ and all $i \in \mathbb{N}$, we have $w, i \models \widehat{\varphi}$ iff $w, i \models \varphi$. ◀

Succinctness. The exponential blow-up predicted by Thm. 8 may seem like a downside at first sight. However, it should be read as the possibility that certain properties, which are definable in LTL, can be defined in sMTL^1 with much shorter formulas. The same

phenomenon is of course known from LTL. It is easy to show that every family of LTL formulas equivalent to the MTL formulas $\varphi_n := \mathtt{F}_{\geq 2^n}\, q$ require size $\mathcal{O}(2^n)$. The proof can be re-used entirely to show that sMTL^1 is also exponentially more succinct than LTL. Note that $|\,\mathtt{F}^1_{\geq 2^n}\, q| = \mathcal{O}(n)$.

▶ **Theorem 9.** *There is a family of LTL-definable languages $(L_n)_{n\geq 1}$ over a singleton $\mathcal{P}$ such that each L_n is expressible in sMTL^1 by a formula of size $\mathcal{O}(n)$ but every family $(\varphi_n)_{n\geq 1}$ of LTL formulas with $L(\varphi_n) = L_n$ is such that $|\varphi_n| = \Omega(2^n)$.*

Proof. Consider the sMTL^1 formulas $\varphi_n := \mathtt{F}^1_{\geq 2^n}\, q$ for $n \geq 1$. Clearly, $|\varphi_n| = \mathcal{O}(n)$. By a standard induction on the structure of LTL formulas we can show that formulas of size $< 2^n$ cannot distinguish between the two words $w_n = \emptyset^{2^n-1}\{q\}\emptyset^\omega$ and $w'_n = \emptyset^{2^n}\{q\}\emptyset^\omega$. Since $w_n \models \varphi_n$ but $w'_n \not\models \varphi_n$ for all $n \geq 1$, we get that any presumed LTL-formula equivalent to φ_n needs to have size 2^n at least. ◀

4 An Automata-Theoretic Decision Procedure

We give an automata-theoretic decision procedure for MTL^1. Decidability of its satisfiability problem is not a surprise in the light of Thm. 8, stating that sMTL^1– and therefore also MTL^1 according to Thm. 7 – can be translated into LTL at an exponential blow-up. This is unavoidable according to Thm. 9. Then it is perhaps rather surprising that the complexity of MTL^1 is asymptotically no worse than that of LTL. We give an upper bound of PSpace, based on a translation into Streett automata. According to Thm. 8, it suffices to do so for sMTL^1.

Temporal formulas and their unfoldings. For convenience, we work with sMTL^1 formulas in *negation normal form* (NNF), i.e. those that are built from literals $q, \neg q$ using the Boolean operators $\wedge, \vee$ and the temporal operators $\mathtt{X}, \mathtt{U}^1$ and $\mathtt{R}^1$ where $\varphi\, \mathtt{R}^1_{\sim k}\, \psi := \neg(\neg\varphi\, \mathtt{U}^1_{\sim k}\, \neg\psi)$. The following is a standard observation about the ability to push negations inwards in formulas to obtain NNF.

▶ **Lemma 10.** *For every sMTL^1 formula φ there is an sMTL^1 formula $\overline{\varphi}$ in NNF of size $\mathcal{O}(|\varphi|)$ such that $\overline{\varphi} \equiv \varphi$.*

The construction of a Streett automaton recognising $L(\varphi)$ for some sMTL^1 formula φ in NNF then follows the same principles as the standard construction of a Büchi automaton for an LTL formula. Temporal operators are typically handled by unfolding, not just in automata-theoretic decision procedures. The term denotes the two equivalences

$$\varphi\, \mathtt{U}\, \psi \equiv \psi \vee (\varphi \wedge \mathtt{X}(\varphi\, \mathtt{U}\, \psi)) \quad \text{and} \quad \varphi\, \mathtt{R}\, \psi \equiv \psi \wedge (\varphi \vee \mathtt{X}(\varphi\, \mathtt{R}\, \psi))\ .$$

These can be extended to the temporal operators in sMTL^1 as follows. The proof is just by close inspection of the semantics of MTL^1.

▶ **Lemma 11.** *Let $\varphi, \psi \in$ sMTL^1, $k \geq 0$. We have*

$$\begin{aligned}
\varphi\, \mathtt{U}^1_{\leq 0}\, \psi &\equiv \psi & \qquad \varphi\, \mathtt{R}^1_{\leq 0}\, \psi &\equiv \psi \\
\varphi\, \mathtt{U}^1_{\geq 0}\, \psi &\equiv \psi \vee (\varphi \wedge \mathtt{X}(\varphi\, \mathtt{U}^1_{\geq 0}\, \psi)) & \varphi\, \mathtt{R}^1_{\geq 0}\, \psi &\equiv \psi \wedge (\varphi \vee \mathtt{X}(\varphi\, \mathtt{R}^1_{\geq 0}\, \psi)) \\
\varphi\, \mathtt{U}^1_{\leq k+1}\, \psi &\equiv \psi \vee (\varphi \wedge \mathtt{X}(\varphi\, \mathtt{U}^1_{\leq k}\, \psi)) & \varphi\, \mathtt{R}^1_{\leq k+1}\, \psi &\equiv \psi \wedge (\varphi \vee \mathtt{X}(\varphi\, \mathtt{R}^1_{\leq k}\, \psi)) \\
\varphi\, \mathtt{U}^1_{\geq k+1}\, \psi &\equiv \neg\psi \wedge \varphi \wedge \mathtt{X}(\varphi\, \mathtt{U}^1_{\geq k}\, \psi) & \varphi\, \mathtt{R}^1_{\geq k+1}\, \psi &\equiv \neg\psi \vee \varphi \vee \mathtt{X}(\varphi\, \mathtt{R}^1_{\geq k}\, \psi)
\end{aligned}$$

A formula χ of the form on one of the left-hand sides in these equations is called a *temporal formula* and we use $unf(\chi)$ to denote the corresponding right-hand side. Temporal formulas of the form $\varphi\,\mathtt{U}^1_{\geq k}\,\psi$, i.e. Until-formulas whose metric parameter is a half-open interval to the right, are called *critical*.

Later on we will need a second observation about temporal formulas. The proof is straightforward from an inspection of the semantics of MTL^1. The lemma already holds for MTL in fact.

▶ **Lemma 12.** *Let φ, ψ be sMTL^1 formulas and $k, \ell \in \mathbb{N}$ such that $k \leq \ell$. Then we have*

$$\models (\varphi\,\mathtt{U}^1_{\leq k}\,\psi) \to (\varphi\,\mathtt{U}^1_{\leq \ell}\,\psi) \qquad \models (\varphi\,\mathtt{R}^1_{\leq \ell}\,\psi) \to (\varphi\,\mathtt{R}^1_{\leq k}\,\psi)$$
$$\models (\varphi\,\mathtt{R}^1_{\geq k}\,\psi) \to (\varphi\,\mathtt{R}^1_{\geq \ell}\,\psi) \qquad \models (\varphi\,\mathtt{U}^1_{\geq \ell}\,\psi) \to (\varphi\,\mathtt{U}^1_{\geq k}\,\psi)$$

The Fischer-Ladner closure of a formula χ is a collection of all subformulas and perhaps others derived from them that may play a role in determining the truth of χ at some point in a word.

▶ **Definition 13.** *Let $\chi \in$ sMTL^1 be in NNF. Its* Fischer-Ladner closure *is the least set $FL(\chi)$ that contains χ and is closed under the following operations.*

- *If $\varphi \wedge \psi \in FL(\chi)$ or $\varphi \vee \psi \in FL(\chi)$ then $\{\varphi, \psi\} \subseteq FL(\varphi)$.*
- *If $\mathtt{X}\,\varphi \in FL(\chi)$ then $\varphi \in FL(\chi)$.*
- *If $\varphi \in FL(\chi)$ for a temporal φ, then $unf(\varphi) \in FL(\chi)$.*

The key concept in an automata construction for sMTL^1 formulas is that of a Hintikka set – a set of formulas that is closed under propositional logic consequence. We need to refine the standard definition slightly in order to obtain the intended complexity bound in the end.

▶ **Definition 14.** Let $\chi \in$ sMTL^1 be in NNF. A $\Phi \subseteq FL(\chi)$ is called a *Hintikka set* for χ if it satisfies the following conditions.

- If $\varphi \wedge \psi \in \Phi$ then $\{\varphi, \psi\} \subseteq \Phi$.
- If $\varphi \vee \psi \in \Phi$ then $\{\varphi, \psi\} \cap \Phi \neq \emptyset$.
- If $\varphi \in \Phi$ for a temporal φ then $unf(\varphi) \in \Phi$.

Φ is called *(propositionally) consistent* if there is no $q \in \mathcal{P}$ such that $\{q, \neg q\} \subseteq \Phi$. It is called *lean* if for all φ, ψ there is at most one temporal formula $\varphi\,\mathtt{U}^1_{\leq k}\,\psi \in \Phi$ for any $k \geq 0$, and likewise for temporal formulas of the three other forms with operators $\mathtt{U}^1_{\geq k}$, $\mathtt{R}^1_{\leq k}$ and $\mathtt{R}^1_{\geq k}$. We write $\mathcal{H}(\chi)$, resp. $\mathcal{H}_{\mathsf{ln}}(\chi)$ for the set of all lean, propositionally consistent Hintikka sets for χ, resp. those that are additionally lean.

While the Fischer-Ladner closure of an sMTL^1 formula χ is generally exponential in $|\chi|$ and there are, thus, doubly exponentially many Hintikka sets, there are only singly exponentially many lean Hintikka sets.

▶ **Lemma 15.** *Let $\chi \in$ sMTL^1 be in NNF. Then $|\mathcal{H}_{\mathsf{ln}}(\chi)| = 2^{\mathcal{O}(|\chi|^2)}$.*

Proof. Note that there are at most $\mathcal{O}(|\chi|)$ many formula schemes, i.e. members of $FL(\chi)$ modulo concrete metric parameters. A lean set can then be seen as a mapping for each such formula scheme to a value in $\{\bot, 0, \ldots, m(\chi)\}$, indicating (non-)inclusion in the set and giving a concrete parameter value for a temporal formula. Since $m(\chi) \in 2^{\mathcal{O}(|\chi|)}$ due to binary encoding, there are at most $(2^{\mathcal{O}(|\chi|)})^{\mathcal{O}(|\chi|)} = 2^{\mathcal{O}(|\chi|^2)}$ many lean (Hintikka) sets. ◀

Streett automata for sMTL^1 formulas. We are now in a position to define a Streett automaton of singly exponential size that recognises exactly the models of an sMTL^1 formula in NNF.

Fix an sMTL^1 formula χ in NNF over some finite $\mathcal{P}$. We will need three constructions on (lean) Hintikka sets for χ. The first one collects all literals in a lean Hintikka set.

$$Now(\Phi) \;:=\; \bigwedge_{q\in\Phi} q \wedge \bigwedge_{\neg q\in\Phi} \neg q$$

It can be seen as the extraction of a propositional formula determining the letter at a position in a word where all formulas in Φ are supposed to be true. On the other hand,

$$Nxts(\Phi) \;:=\; \{ln(\Psi) \mid \Psi \in \mathcal{H}(\chi) \text{ and } \forall\, \mathtt{X}\,\psi \in \Phi : \psi \in \Psi\}$$

collects all "leanifications" of Hintikka sets that are potential successors to Φ in the sense that they contain all formula ψ which Φ needs to be true at the next position. The leanification $ln(\Psi)$ of Ψ is obtained by successively replacing temporal formulas as follows until no further steps are applicable.

- If $\{\varphi\,\mathtt{U}^1_{\leq k}\,\psi, \varphi\,\mathtt{U}^1_{\leq \ell}\,\psi\} \subseteq \Psi$ or $\{\varphi\,\mathtt{R}^1_{\geq k}\,\psi, \varphi\,\mathtt{R}^1_{\geq \ell}\,\psi\} \subseteq \Psi$ for some φ, ψ and $k < \ell$ then replace all occurrences of the latter that do not occur under the scope of a $\mathtt{X}$-operator by the former.
- If $\{\varphi\,\mathtt{U}^1_{\geq k}\,\psi, \varphi\,\mathtt{U}^1_{\geq \ell}\,\psi\} \subseteq \Psi$ or $\{\varphi\,\mathtt{R}^1_{\leq k}\,\psi, \varphi\,\mathtt{R}^1_{\leq \ell}\,\psi\} \subseteq \Psi$ for some φ, ψ and $k < \ell$ then replace all occurrences of the former that do not occur under the scope of a $\mathtt{X}$-operator by the latter.

It should be clear that $Nxts(\Phi)$ is indeed a set of lean Hintikka sets for χ.

The important observation about the leanification process is the preservation of the semantics in a strong sense.

▶ **Lemma 16.** *Let $\chi \in sMTL^1$ be in NNF and $\Phi \in \mathcal{H}(\chi)$. For every $\varphi \in \Phi$ there is $\varphi' \in ln(\Phi)$ such that φ' differs from φ only in the values of metric parameters and $\models \varphi' \to \varphi$.*

Proof. This is proved in a straight-forward induction on the structure of φ. The only non-trivial cases are those of temporal formulas. These are covered by Lemma 12. Since leanification also replaces subformulas, we also need the fact that χ (and all its subformulas) are given in NNF. Hence, if $\models \varphi' \to \varphi$ then also $\models \psi \to \psi[\varphi/\varphi']$ for any ψ, i.e. all formulas are monotonic. ◀

The acceptance condition of the NSA $\mathcal{A}_\chi$ is determined by the set of all critical temporal formulas that can occur in lean Hintikka sets for χ. However, unfolding decreases interval bounds by one, so the number of critical temporal formulas in $FL(\chi)$ is exponential: if χ contains the subformula $\varphi\,\mathtt{U}^1_{\geq k}\,\psi$ for some $k \in \mathbb{N}$, then $FL(\chi)$ contains $\varphi\,\mathtt{U}^1_{\geq h}\,\psi$ for all $h \leq k$, i.e. exponentially many in $|\chi|$ because of binary encodings of interval bounds.

The exact interval bounds are irrelevant for the acceptance condition. It is only used to ensure that no critical temporal formula $\varphi\,\mathtt{U}^1_{\geq k}\,\psi$ gets unfolded forever without ever satisfying its right argument ψ. We introduce the notion of a critical formula *scheme* $\varphi\,\mathtt{U}^1_{\geq *}\,\psi$ and write $(\varphi\,\mathtt{U}^1_{\geq *}\,\psi) \in \Phi$ if there is some $k \in \mathbb{N}$ such that $(\varphi\,\mathtt{U}^1_{\geq k}\,\psi) \in \Phi$.

Note that there is no need to treat the other three kinds of temporal formulas in the same way. A temporal Release-formula is a greatest fixpoint, and unfolding it infinitely often is a legitimate way of determining its truth. A non-critical, temporal Until-formula of the form $\varphi\,\mathtt{U}^1_{\leq k}\,\psi$ cannot get unfolded infinitely often because each unfolding step decreases the metric parameter in it until it eventually becomes 0, and the formula is replaced by ψ anyway. Note

that this is not the case for critical Until-formulas. Unfolding them still decreases the metric parameter. However, the leanification process can increase it again, whereas leanification for non-critical Until-formulas can only decrease it further.

Let $\{\gamma_1, \dots, \gamma_n\}$ be the set of all schemes of critical temporal Until-formulas in $FL(\chi)$, i.e. $\gamma_i = \alpha_i \,\mathtt{U}^1_{\geq *}\, \beta_i$ for some α_i, β_i. We define the NSA $\mathcal{A}_\chi$ as $(\mathcal{H}_{\mathsf{ln}}(\chi), \mathcal{P}, I, \delta, \mathcal{F})$ where

- $I := \{\Phi \in \mathcal{H}_{\mathsf{ln}}(\chi) \mid \chi \in \Phi\}$,
- $\delta := \{(\Phi, \mathit{Now}(\Phi), \Psi) \mid \Psi \in \mathit{Nxts}(\Phi)\}$,
- $\mathcal{F} := \{(F_1, G_1), \dots, (F_n, G_n)\}$ with $F_i := \{\Phi \mid \gamma_i \in \Phi\}$ and $G_i := \{\Phi \mid \beta_i \in \Phi\}$. Note that β_i may contain other temporal formulas, so β_i is to be understood as a scheme potentially, and $\beta_i \in \Phi$ means that some formula deviating from β_i in metric parameters only is contained in Φ.

The next two lemmas are devoted to the soundness and completeness of the construction.

▶ **Lemma 17.** *Let $\chi \in sMTL^1$ over $\mathcal{P}$ be in NNF and $\mathcal{A}_\chi$ as above. Then $L(\mathcal{A}_\chi) \subseteq L(\chi)$.*

Proof. Let $w = a_0 a_1 \dots (2^{\mathcal{P}})^\omega$ be such that there is an accepting run $\rho = \Phi_0, \Phi_1, \dots$ of $\mathcal{A}_\chi$ on w. By the construction of $\mathcal{A}_\chi$ we have (I) $\chi \in \Phi_0$ and, for all $i \geq 0$, (II) $a_i \models \mathit{Now}(\Phi_i)$ and (III) there is $\Psi_{i+1} \in \mathcal{H}(\chi)$ such that $\psi \in \Psi_{i+1}$ for all $\mathtt{X}\,\psi \in \Phi_i$ and $\Phi_{i+1} = \mathit{ln}(\Psi_{i+1}) \in \mathit{Nxts}(\Phi_i)$.

We show by induction on the structure of formulas φ that for all $i \in \mathbb{N}$ and all $\varphi \in \Phi_i$ we have $w, i \models \varphi$. For literals q or $\neg q$ this follows immediately from (II). For conjunctions and disjunctions this follows by the hypothesis for both conjuncts, resp. one disjunct and the fact that each Φ_i is a lean Hintikka set that behaves like a Hintikka set in this case, i.e. it contains both conjuncts of a conjunction etc. This is the case because leanification replaces either none or all occurrences that are not under the scope of a $\mathtt{X}$-operator. Hence, a replacement takes place in a conjunction iff it takes place in both conjuncts etc.

Suppose φ is of the form $\mathtt{X}\,\psi$. Because of (III), there is a $\Psi_{i+1} \in \mathcal{H}(\chi)$ with $\psi \in \Psi_{i+1}$ with $\Phi_{i+1} = \mathit{ln}(\Psi_{i+1})$. According to Lemma 16, there is $\psi' \in \Phi_{i+1}$ that is structurally not greater than ψ. Hence, we can apply the induction hypothesis to it and obtain $w, i+1 \models \psi'$. According to Lemma 16, we then also have $w, i+1 \models \psi$ and therefore $w, i \models \varphi$.

Suppose φ is of the form $\psi_1 \,\mathtt{U}^1_{\leq k}\, \psi_2$. By inspection of the unfolding rule (Lemma 11) and successive applications of the principles (II) and (III) with the same kind of reasoning using Lemma 16, we get some $k' \leq k$ and a sequence $\Psi_{i+1}, \dots, \Psi_{i+k'}$ of Hintikka sets such that $\Phi_{i+j} = \mathit{ln}(\Psi_{i+j})$, $\psi' \in \Psi_{i+k'}$ for some ψ' with $\models \psi'_2 \to \psi_2$ and $\psi''_h \in \Phi_{i+h}$ for some $\psi''_0, \dots, \psi''_{k'-1}$ such that $\models \psi''_h \to \psi_1$. Applying the induction hypothesis to ψ'_2 at position $i+k'$ and for $\psi''_0, \dots, \psi''_{k'-1}$ at positions $i, \dots, i+k'-1$ shows that these are satisfied at the respective positions in w. Lemma 16 then yields $w, i+k' \models \psi_2$ and $w, i+h \models \psi_1$ for all $h = 0, \dots, k'-1$. Thus, $w, i \models \varphi$.

Suppose φ is of the form $\psi_1 \,\mathtt{U}^1_{\geq k}\, \psi_2$. Note that it is a critical temporal formula in this case. We can apply the same reasoning as in the previous case using Lemmas 11 and 16. However, here the leanification process may replace metric parameters by larger ones. Hence, this alone does not guarantee the existence of a $k' \geq k$ such that $\psi_2 \in \Phi_{i+k'}$. This is where $\mathcal{A}_\chi$'s acceptance condition comes into place. Since ρ is accepting, it must either contain finitely many lean Hintikka sets containing φ or infinitely many containing ψ_2. The latter case immediately implies the existence of such a $k' \geq k$. The former case does so, too, by inspection of Lemma 11 and the construction of Hintikka sets. It is only possible to have (the scheme) φ finitely often only when some $\Phi_{i+k'}$ contains $\psi_1 \,\mathtt{U}^1_{\geq k}\, \psi_2$ and therefore also ψ_2. The rest of this case is handled as the previous one.

The remaining two cases of temporal Release-formulas are also handled in a way that is analogous to those of the Until-formulas.

At last, (I) says that $\chi \in \Phi_0$. By the reasoning above we then have $w, 0 \models \chi$, i.e. $w \in L(\chi)$ which completes the claim. ◀

▶ **Lemma 18.** *Let $\chi \in sMTL^1$ over $\mathcal{P}$ be in NNF and $\mathcal{A}_\chi$ as above. Then $L(\chi) \subseteq L(\mathcal{A}_\chi)$.*

Proof. Suppose $w = a_0 a_1 \ldots \in L(\chi)$. We need to construct an accepting run $\Phi_0, \Phi_1, \ldots$ of $\mathcal{A}_\chi$ on w. To this end, we construct a sequence $\Phi'_0, \Phi'_1, \ldots$ via $\Phi'_i := \{\varphi \in FL(\chi) \mid w, i \models \varphi\}$. It is not hard to see that each Φ'_i is indeed a propositionally consistent Hintikka set. Moreover, if $\mathtt{X}\varphi \in \Phi_i$ then $\varphi \in \Phi_{i+1}$. This therefore determines a sequence of lean Hintikka sets by leanification: $\Phi_i := ln(\Phi'_i)$ for all $i \geq 0$, forming a run $\rho = \Phi_0, \Phi_1, \ldots$.

It remains to be seen that it is indeed an accepting run. Take some critical Until-formula scheme $\gamma = \alpha \,\mathtt{U}^1_{\geq *}\, \beta$ and suppose that ρ contains an infinite subsequence $\Phi_{i_1}, \Phi_{i_2}, \ldots$ with $\gamma_j := (\alpha \,\mathtt{U}^1_{\geq k_j}\, \beta) \in \Phi_{i_j}$ for all $j \geq 1$. By construction, we have $w, i_j \models \gamma_j$. By the semantics of MTL^1 there are $k'_1, k'_2, \ldots$ with $k'_j \geq k_j$ such that $w, i_j + k'_j \models \beta$ and, again, by construction $\beta \in \Phi'_{i_j + k'_j}$. Since $i_1 < i_2 < \ldots$, the set $\{i_j + k'_j \mid j \geq 1\}$ is infinite. Hence, the sequence $\Phi'_0, \Phi'_1, \ldots$ has an infinite subsequence in which every Hintikka set contains β. Then the run ρ has an infinite subsequence in which every lean Hintikka set either contains β itself or an instantiation of the scheme β. In any case, the run ρ satisfies the Streett pair associated with γ. Since this is the case for any critical γ, ρ is indeed accepting and so we have $w \in L(\mathcal{A}_\chi)$. ◀

Putting all of the above together we obtain that MTL^1 is not just decidable but that its satisfiability problem is no worse than that of ordinary LTL (and therefore exponentially easier than that of MTL), even though there is an exponential succinctness gap between MTL^1 and LTL.

▶ **Theorem 19.** *SAT(MTL^1) is PSpace-complete.*

Proof. The lower bound is straightforwardly inherited from LTL. For the upper bound, note that every MTL^1 formula χ can be translated into an equivalent sMTL^1 formula in NNF at a linear blow-up only (Thm. 7 and Lemma 10). This can in turn be translated into an equivalent NSA (Lemmas 17 and 18) of exponential size in $|\chi|$ (Lemma 15). Language equivalence entails particularly that its language is non-empty iff χ is satisfiable. Non-emptiness for NSA can be decided in NLogSpace (Thm. 4) which is NPSpace measured in the size of χ. Savitch's Theorem [15] then gives a PSpace upper bound. ◀

5 Conclusion

We have investigated the expressiveness and computational complexity of MTL^1, a variant of Metric Linear-Time Temporal logic MTL in wich the metric parameters do not constrain the occurrence of some event but their first occurrence (in an Until formula). The resulting logic is still exponentially more succinct than LTL. Unlike full MTL whose satisfiability problem is ExpSpace-complete, we obtained a PSpace upper bound for MTL^1 by the construction of equivalent Streett automata of exponential size.

The main motivation for the study of this logic is given by links to State-Space Models in machine learning. Further work will elaborate on the connections between these formalisms. On the side of temporal logics, there is obvious further work in terms generalisations of the strict first-time semantics to a strict n-th time semantics, constraining further moments in which an Until-formula gets satisfied. We suspect that for every fixed n, the resulting logic MTL^n remains PSpace-complete.

There are also obvious connections to Counting LTL [11], a variant of LTL with a counting operator. The first-time semantics is clearly expressible using counting operators by stating that the number of positions beforehand is zero. The relative succinctness between the two formalisms remains to be investigated, as is the case for MTL and MTL[1].

References

1 R. Alur and T. Henzinger. Real-time logics: Complexity and expressiveness. *Information and Computation*, 104(1):35–77, 1993. doi:10.1006/INCO.1993.1025.

2 S. Arora and B. Barak. *Computational Complexity: A Modern Approach*. Cambridge Univ. Press, 2006.

3 B. Banieqbal and H. Barringer. Temporal logic with fixed points. In *Proc. Coll. on Temporal Logic in Specification*, volume 398 of *LNCS*, pages 62–73. Springer, 1989. doi:10.1007/3-540-51803-7_22.

4 S. Demri, V. Goranko, and M. Lange. *Temporal Logics in Computer Science*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2016.

5 C. Eisner and D. Fisman. *A Practical Introduction to PSL*. Series on Integrated Circuits and Systems. Springer, 2006. doi:10.1007/978-0-387-36123-9.

6 D. Gabbay, A. Pnueli, S. Shelah, and J. Stavi. On the temporal analysis of fairness. In *Proc. 7th Symp. on Principles of Programming Languages, POPL'80*, pages 163–173. ACM, 1980. doi:10.1145/567446.567462.

7 G. De Giacomo and M. Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *Proc. 23rd Int. Joint Conf. on A.I., IJCAI'13*, pages 854–860. IJCAI/AAAI, 2013. URL: http://www.aaai.org/ocs/index.php/IJCAI/IJCAI13/paper/view/6997.

8 M. Rauch Henzinger and J. A. Telle. Faster algorithms for the nonemptiness of streett automata and for communication protocol pruning. In *Proc. 5th Scand. Workshop on Algorithm Theory, SWAT'96*, volume 1097 of *LNCS*, pages 16–27. Springer, 1996. doi:10.1007/3-540-61422-2_117.

9 Ron Koymans. Specifying real-time properties with metric temporal logic. *Real Time Syst.*, 2(4):255–299, 1990. doi:10.1007/BF01995674.

10 F. Laroussinie, N. Markey, and Ph. Schnoebelen. Efficient timed model checking for discrete-time systems. *Theoretical Computer Science*, 353(1):249–271, 2006. doi:10.1016/j.tcs.2005.11.020.

11 F. Laroussinie, A. Meyer, and E. Petonnet. Counting LTL. In *Proc. 17th Int. Symp. on Temporal Representation and Reasoning, TIME'10*, pages 51–58. IEEE, 2010. doi:10.1109/TIME.2010.20.

12 T. Latvala and K. Heljanko. Coping with strong fairness. *Fundam. Informaticae*, 43(1-4):175–193, 2000. doi:10.3233/FI-2000-43123409.

13 W. Merrill, J. Petty, and A. Sabharwal. The illusion of state in state-space models. In *Proc. 41st Int. Conf. on Machine Learning, ICML'24*, volume 235, pages 35492–35506. JMLR.org, 2024.

14 A. Pnueli. The temporal logic of programs. In *Proc. 18th Symp. on Foundations of Computer Science, FOCS'77*, pages 46–57. IEEE, 1977. doi:10.1109/SFCS.1977.32.

15 W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4:177–192, 1970. doi:10.1016/S0022-0000(70)80006-X.

16 A. P. Sistla. *Theoretical Issues in the Design of Distributed and Concurrent Systems*. PhD thesis, Harvard Univ., Cambridge, MA, 1983.

17 A. P. Sistla and E. M. Clarke. The complexity of propositional linear temporal logics. *Journal of the ACM*, 32(3):733–749, 1985. doi:10.1145/3828.3837.

18 R. S. Streett. Propositional dynamic logic of looping and converse is elementarily decidable. *Information and Control*, 54(1/2):121–141, 1982. doi:10.1016/S0019-9958(82)91258-X.

19 M. Y. Vardi. A temporal fixpoint calculus. In *Proc. Conf. on Principles of Programming Languages, POPL'88*, pages 250–259. ACM, 1988. doi:10.1145/73560.73582.

20 M. Y. Vardi. *An Automata-Theoretic Approach to Linear Temporal Logic*, volume 1043 of *LNCS*, pages 238–266. Springer, New York, NY, USA, 1996. doi:10.1007/3-540-60915-6_6.

21 M. Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information and Computation*, 115(1):1–37, 1994. doi:10.1006/INCO.1994.1092.