

The Complexity of Clausal Fragments of LTL

Alessandro Artale¹, Roman Kontchakov², Vladislav Ryzhikov¹,
and Michael Zakharyashev²

¹ KRDB Research Centre
Free University of Bozen-Bolzano
I-39100 Bolzano, Italy
{artale,ryzhikov}@inf.unibz.it

² Department of Computer Science and Information Systems
Birkbeck, University of London
London WC1E 7HX, UK
{roman,michael}@dcs.bbk.ac.uk

Abstract. We introduce and investigate a number of fragments of propositional temporal logic LTL over the flow of time $(\mathbb{Z}, <)$. The fragments are defined in terms of the available temporal operators and the structure of the clausal normal form of the temporal formulas. We determine the computational complexity of the satisfiability problem for each of the fragments, which ranges from NLOGSPACE to PTIME, NP and PSPACE.

1 Introduction

We consider the (PSPACE-complete) propositional temporal logic LTL over the flow of time $(\mathbb{Z}, <)$. Our aim is to investigate how the computational complexity of the satisfiability problem for LTL-formulas depends on the form of their clausal representation and the available temporal operators.

Sistla and Clarke [26] showed that satisfiability of LTL-formulas with all standard operators ('next-time', 'always in the future', 'eventually' and 'until') is PSPACE-complete; see also [18,19]. Ono and Nakamura [22] proved that for formulas with only 'always in the future' and 'eventually' the satisfiability problem becomes NP-complete. Since then a number of fragments of LTL of different complexity have been identified. For example, Chen and Lin [10] observed that the complexity does not change if we restrict attention to temporal Horn formulas. Demri and Schnoebelen [12] determined the complexity of fragments that depend on three parameters: the available temporal operators, the number of nested temporal operators, and the number of propositional variables in formulas. Markey [21] analysed fragments defined by the allowed set of temporal operators, their nesting and the use of negation. Dixon *et al.* [13] introduced a XOR fragment of LTL and showed its tractability. Bauland *et al.* [7] systematically investigated the complexity of fragments given by both temporal operators and Boolean connectives (using Post's lattice of sets of Boolean functions).

In this paper, we classify temporal formulas according to their clausal normal form. Recall [14] that any LTL-formula over $(\mathbb{N}, <)$ can be transformed into an equisatisfiable formula in the so-called *separated normal form* that consists of

Table 1. The complexity of clausal fragments of LTL

temporal operators	$\boxtimes, \square_F, \square_P, \circ_F, \circ_P$	$\boxtimes, \square_F, \square_P$	\boxtimes
α	$\text{LTL}_{\alpha}^{\square, \circ}$	$\text{LTL}_{\alpha}^{\square}$	$\text{LTL}_{\alpha}^{\boxtimes}$
<i>bool</i>	PSPACE (\leq [26])	NP (\leq [22])	NP
<i>horn</i>	PSPACE (\geq [10])	PSPACE (\leq Th. 3)	PSPACE
<i>krom</i>	NP (\leq Th. 1)	NP (\geq Th. 5)	NLOGSPACE
<i>core</i>	NP (\geq Th. 2)	NLOGSPACE (\leq Th. 4)	NLOGSPACE

initial clauses (setting conditions at moment 0), step clauses (defining transitions between consecutive states), and eventuality clauses (defining the states that must be reached infinitely often). Our clausal normal form is a slight generalisation of the separated normal form. The main building blocks are *positive temporal literals* λ given by the following grammar:

$$\lambda ::= \perp \mid p \mid \circ_F \lambda \mid \circ_P \lambda \mid \square_F \lambda \mid \square_P \lambda \mid \boxtimes \lambda, \quad (1)$$

where p is a propositional variable, \circ_F and \circ_P are the next- and previous-time operators, and \square_F , \square_P , \boxtimes are the operators ‘always in the future,’ ‘always in the past’ and ‘always.’ We say that a temporal formula φ is in *clausal normal form* if

$$\varphi ::= \lambda \mid \neg \lambda \mid \boxtimes (\neg \lambda_1 \vee \dots \vee \neg \lambda_n \vee \lambda_{n+1} \vee \dots \vee \lambda_{n+m}) \mid \varphi_1 \wedge \varphi_2. \quad (2)$$

Conjunctions of positive and *negative* ($\neg \lambda$) literals can be thought of as initial clauses, while conjunctions of \boxtimes -formulas generalise both step and eventuality clauses of the separated normal form. Similarly to [15] one can show that any LTL-formula over $(\mathbb{Z}, <)$ is equisatisfiable to a formula in clausal normal form.

We consider twelve fragments of LTL that will be denoted by $\text{LTL}_{\alpha}^{\square, \circ}$, $\text{LTL}_{\alpha}^{\square}$ and $\text{LTL}_{\alpha}^{\boxtimes}$, for $\alpha \in \{\text{bool}, \text{horn}, \text{krom}, \text{core}\}$. The superscript in the language name indicates the temporal operators that can be used in its positive literals. Thus, $\text{LTL}_{\alpha}^{\square, \circ}$ uses all types of positive literals, $\text{LTL}_{\alpha}^{\square}$ can only use the \square -operators:

$$\lambda ::= \perp \mid p \mid \square_F \lambda \mid \square_P \lambda \mid \boxtimes \lambda,$$

and $\text{LTL}_{\alpha}^{\boxtimes}$ only the \boxtimes -operator:

$$\lambda ::= \perp \mid p \mid \boxtimes \lambda.$$

The subscript α in the language name refers to the form of the clauses

$$\neg \lambda_1 \vee \dots \vee \neg \lambda_n \vee \lambda_{n+1} \vee \dots \vee \lambda_{n+m} \quad (3)$$

($m, n \geq 0$) that can be used in the formulas φ :

- *bool*-clauses are arbitrary clauses of the form (3),
- *horn*-clauses have at most one positive literal (that is, $m \leq 1$),

- *krom*-clauses are binary (that is, $n + m \leq 2$),
- *core*-clauses are binary with at most one positive literal ($n + m \leq 2$, $m \leq 1$).

The tight complexity bounds in Table 1 show how the complexity of the satisfiability problem for LTL-formulas depends on the form of clauses and the available temporal operators. The PSPACE upper bound for $\text{LTL}_{bool}^{\square, \circ}$ is well-known [18,26,24,25]; the matching lower bound can be obtained already for $\text{LTL}_{horn}^{\square, \circ}$ without \square_F and \square_P by a standard encoding of deterministic Turing machines with polynomial tape [10]. The NP upper bound for $\text{LTL}_{horn}^{\square}$ is also well-known [22], and the PTIME and NLOGSPACE lower bounds for $\text{LTL}_{horn}^{\boxplus}$ and $\text{LTL}_{core}^{\boxplus}$ coincide with the complexity of the respective non-temporal languages. The upper bounds for the $\text{LTL}_{\alpha}^{\boxplus}$ fragments can be obtained by embedding into the underlying propositional fragments; see the full paper [6] for details.

The main contributions of this paper are the remaining complexity results in Table 1. The complexity of the $\text{LTL}_{\alpha}^{\square}$ fragments matches the complexity of the underlying non-temporal fragments except for the Krom case, where we can use the clauses $\neg p \vee \neg \square_F q$ and $q \vee r$ to say that $p \rightarrow \diamond_F r$ (if p then eventually r), which allows one to encode 3-colourability and results in NP-hardness. It is known that the addition of the operators \circ_F and \circ_P to the language with \square_F and \square_P usually increases the complexity (note that the proofs of the lower bounds for the $\text{LTL}_{\alpha}^{\square, \circ}$ fragments require only \boxplus and \circ_F). It is rather surprising that this does not happen in the case of the Krom fragment, while the complexity of the corresponding core fragment jumps from NLOGSPACE to NP.

We prove the upper bounds using two different techniques. The existence of models for $\text{LTL}_{krom}^{\square, \circ}$ -formulas is checked in Section 3 by guessing a small number of types and exponentially large distances between them (given in binary) and then using unary automata (and the induced arithmetic progressions) to verify correctness of the guess in polynomial time. In Section 4.1, we design a calculus for $\text{LTL}_{core}^{\square}$ in which derivations can be thought of as paths in a graph over the propositions labelled by moments of time. Thus, the existence of such derivations is essentially the graph reachability problem and can be solved in NLOGSPACE.

2 The Clausal Normal Form for LTL

The *propositional linear-time temporal logic* LTL (see, e.g., [16,17] and references therein) we consider in this paper is interpreted over the flow of time $(\mathbb{Z}, <)$. LTL-formulas are built from propositional variables p_0, p_1, \dots , propositional constants \top and \perp , the Boolean connectives $\wedge, \vee, \rightarrow$ and \neg , and two binary temporal operators \mathcal{S} (‘since’) and \mathcal{U} (‘until’), which are assumed to be ‘strict.’ So, the other temporal operators mentioned in the introduction can be defined via \mathcal{S} and \mathcal{U} as follows:

$$\begin{array}{llll}
 \circ_F \varphi = \perp \mathcal{U} \varphi, & \diamond_F \varphi = \top \mathcal{U} \varphi, & \square_F \varphi = \neg \diamond_F \neg \varphi, & \boxtimes \varphi = \diamond_P \diamond_F \varphi, \\
 \circ_P \varphi = \perp \mathcal{S} \varphi, & \diamond_P \varphi = \top \mathcal{S} \varphi, & \square_P \varphi = \neg \diamond_P \neg \varphi, & \boxplus \varphi = \square_P \square_F \varphi.
 \end{array}$$

A *temporal interpretation*, \mathfrak{M} , defines a truth-relation between moments of time $n \in \mathbb{Z}$ and propositional variables p_i . We write $\mathfrak{M}, n \models p_i$ to indicate that p_i is true at the moment n in the interpretation \mathfrak{M} . This truth-relation is extended to all LTL-formulas as follows (the Booleans are interpreted as expected):

$\mathfrak{M}, n \models \varphi \mathcal{U} \psi$ iff there is $k > n$ with $\mathfrak{M}, k \models \psi$ and $\mathfrak{M}, m \models \varphi$, for $n < m < k$,
 $\mathfrak{M}, n \models \varphi \mathcal{S} \psi$ iff there is $k < n$ with $\mathfrak{M}, k \models \psi$ and $\mathfrak{M}, m \models \varphi$, for $k < m < n$.

An LTL-formula φ is *satisfiable* if there is an interpretation \mathfrak{M} such that $\mathfrak{M}, 0 \models \varphi$; in this case we call \mathfrak{M} a *model* of φ . We denote the length of φ by $|\varphi|$.

Recall that LTL-formulas of the form (2) were said to be in *clausal normal form*, and the class of such formulas was denoted by $\text{LTL}_{bool}^{\square, \circ}$. The clauses (3) will often be represented as $\lambda_1 \wedge \dots \wedge \lambda_n \rightarrow \lambda_{n+1} \vee \dots \vee \lambda_{n+m}$ (where the empty disjunction is \perp and the empty conjunction is \top).

Lemma 1 (clausal normal form). *For every LTL-formula, one can construct an equisatisfiable $\text{LTL}_{bool}^{\square, \circ}$ -formula. The construction requires logarithmic space.*

The proof of this lemma is similar to the proof of [15, Theorem 3.3.1] and uses fixed-point unfolding and renaming [15,23]. For example, we can replace every positive occurrence (that is, an occurrence in the scope of an even number of negations) of $p \mathcal{U} q$ in a given formula φ with a fresh propositional variable r and add the conjuncts $\boxtimes(r \rightarrow \bigcirc_F q \vee \bigcirc_F p)$, $\boxtimes(r \rightarrow \bigcirc_F q \vee \bigcirc_F r)$ and $\boxtimes(r \rightarrow \diamond_F q)$ to φ . The result contains no positive occurrences of $p \mathcal{U} q$ and is equisatisfiable with φ : the first two conjuncts are the fixed-point unfolding ($p \mathcal{U} q \rightarrow \bigcirc_F q \vee (\bigcirc_F p \wedge \bigcirc_F (p \mathcal{U} q))$), while the last conjunct ensures that the fixed-point is eventually reached.

The next lemma allows us to consider an even more restricted classes of formulas. In what follows, we do not distinguish between a set of formulas and the conjunction of its members, and we write $\boxtimes \Phi$ for the conjunction $\bigwedge_{\chi \in \Phi} \boxtimes \chi$.

Lemma 2. *Let \mathcal{L} be one of $\text{LTL}_{\alpha}^{\square, \circ}$, $\text{LTL}_{\alpha}^{\square}$, $\text{LTL}_{\alpha}^{\boxtimes}$, for $\alpha \in \{bool, horn, krom, core\}$. For any \mathcal{L} -formula φ , one can construct, in log-space, an equisatisfiable \mathcal{L} -formula*

$$\Psi \wedge \boxtimes \Phi, \quad (4)$$

where Ψ is a conjunction of propositional variables from Φ , and Φ is a conjunction of clauses of the form (3) containing only \bigcirc_F , \square_P , \square_F for $\text{LTL}_{\alpha}^{\square, \circ}$, only \square_P , \square_F for $\text{LTL}_{\alpha}^{\square}$, and only \boxtimes for $\text{LTL}_{\alpha}^{\boxtimes}$, in which the temporal operators are not nested.

Proof. First, we take a fresh variable p and replace all the conjuncts of the form λ and $\neg \lambda$ in φ by $\boxtimes(\neg p \vee \lambda)$ and $\boxtimes(\neg p \vee \neg \lambda)$, respectively; we set $\Psi = p$. For an $\text{LTL}_{\alpha}^{\square, \circ}$ or $\text{LTL}_{\alpha}^{\square}$ -formula, we replace the temporal literals $\boxtimes \lambda$ with $\square_F \square_P \lambda$. Then, for each $\bigcirc_P \lambda$, we take a fresh variable, denoted $\overline{\bigcirc_P \lambda}$, replace each occurrence of $\bigcirc_P \lambda$ with $\overline{\bigcirc_P \lambda}$ and add the conjuncts $\boxtimes(\bigcirc_F \overline{\bigcirc_P \lambda} \rightarrow \lambda)$ and $\boxtimes(\lambda \rightarrow \bigcirc_F \overline{\bigcirc_P \lambda})$ to the resulting formula. In a similar manner, we use fresh

where $\bigcirc_F^k \Psi$ is the result of attaching k operators \bigcirc_F to each literal in Ψ and

$$\begin{aligned} \Theta_i = & \{p, \overline{\bigcirc_F p} \mid \overline{\bigcirc_F p} \in \Psi_i\} \cup \{\neg \overline{\bigcirc_F p} \mid \neg \overline{\bigcirc_F p} \in \Psi_i\} \cup \\ & \{p, \overline{\bigcirc_F p} \mid \overline{\bigcirc_F p} \in \Psi_{i+1}\} \cup \{\neg \overline{\bigcirc_F p} \mid \neg \overline{\bigcirc_F p} \in \Psi_{i+1}\}. \end{aligned}$$

The intuition behind this lemma is as follows (see Fig. 1). If φ is satisfiable, then it has a model \mathfrak{M} that consists of the initial fragments of models \mathfrak{M}_i of the formulas ψ_i : namely, the types of the moments m_i, \dots, m_{i+1} in \mathfrak{M} coincide with the types of the moments $0, \dots, (m_{i+1} - m_i)$ in \mathfrak{M}_i . By **(B₄)**, we have $\mathfrak{M}, 0 \models \boxtimes \overline{\Phi}$. Then **(B₁)** makes sure that $\mathfrak{M}, 0 \models \Psi$. Conditions **(B₂)** and **(B₃)** guarantee that if $\overline{\bigcirc_F p} \in \Psi_i$ then $p \in \Psi_j$ for all types Ψ_j located to the right of Ψ_i in Fig. 1 and, conversely, if $\overline{\bigcirc_F p} \notin \Psi_i$ then $\neg p \in \Psi_j$, for some Ψ_j to the right of Ψ_i ; and symmetrically for the \bigcirc_F -literals. It follows that $\mathfrak{M}, 0 \models \boxtimes \Phi$.

3 Binary-Clause LTL and Arithmetic Progressions

In this section, we prove NP-completeness of the satisfiability problem for $\text{LTL}_{krom}^{\square, \bigcirc}$ and $\text{LTL}_{core}^{\square, \bigcirc}$. The key ingredient of the proof of the upper bound is an encoding of condition **(B₄)** for *binary clauses* by means of arithmetic progressions (via unary automata). The proof of the lower bound is by reduction of the problem whether a given set of arithmetic progressions covers all the natural numbers.

Let φ be an $\text{LTL}_{krom}^{\square, \bigcirc}$ -formula of the form (4). By Lemma 3, to check satisfiability of φ it suffices to guess $K + 1$ types for $\overline{\Phi}$ and K natural numbers $n_i = m_{i+1} - m_i$, for $0 \leq i < K$, whose binary representation, by **(B₀)**, is polynomial in $|\overline{\Phi}|$. Evidently, **(B₁)**–**(B₃)** can be checked in polynomial time. Our aim now is to show that **(B₄)** can also be verified in polynomial time, which will give a nondeterministic polynomial-time algorithm for checking satisfiability of $\text{LTL}_{krom}^{\square, \bigcirc}$ -formulas.

Theorem 1. *The satisfiability problem for $\text{LTL}_{krom}^{\square, \bigcirc}$ -formulas is in NP.*

Proof. In view of Lemma 2, we write \bigcirc in place of \bigcirc_F . We denote propositional literals (p or $\neg p$) by L and temporal literals (p , $\neg p$, $\bigcirc p$ or $\neg \bigcirc p$) by D . We assume that $\bigcirc \neg p$ is the same as $\neg \bigcirc p$. We use $\psi_1 \models \psi_2$ as a shorthand for ‘ $\mathfrak{M}, 0 \models \psi_2$ whenever $\mathfrak{M}, 0 \models \psi_1$, for any interpretation \mathfrak{M} .’ Thus, the problem is as follows: given a set $\overline{\Phi}$ of binary clauses of the form $D_1 \vee D_2$, types Ψ and Ψ' for $\overline{\Phi}$, a set Θ of propositional literals and a number $n > 0$ (in binary), decide whether

$$\Psi \wedge \bigwedge_{k=1}^{n-1} \bigcirc^k \Theta \wedge \bigcirc^n \Psi' \wedge \boxtimes \overline{\Phi} \quad (5)$$

has a satisfying interpretation. For $0 \leq k \leq n$, we set:

$$\begin{aligned} F_{\overline{\Phi}}^k(\Psi) &= \{L' \mid L \wedge \boxtimes \overline{\Phi} \models \bigcirc^k L', \text{ for } L \in \Psi\}, \\ P_{\overline{\Phi}}^k(\Psi') &= \{L \mid \bigcirc^k L' \wedge \boxtimes \overline{\Phi} \models L, \text{ for } L' \in \Psi'\}. \end{aligned}$$

Lemma 4. *Formula (5) is satisfiable iff the following conditions hold:*

- (**L**₁) $F_{\Phi}^0(\Psi) \subseteq \Psi$, $F_{\Phi}^n(\Psi) \subseteq \Psi'$ and $F_{\Phi}^0(\Psi') \subseteq \Psi'$, $P_{\Phi}^n(\Psi') \subseteq \Psi$;
 (**L**₂) $\neg L \notin F_{\Phi}^k(\Psi)$ and $\neg L \notin P_{\Phi}^{n-k}(\Psi')$, for all $L \in \Theta$ and $0 < k < n$.

Proof. Clearly, if (5) is satisfiable then the above conditions hold. For the converse direction, observe that if $L' \in F_{\Phi}^k(\Psi)$ then, since Φ is a set of binary clauses, there is a sequence of \circ -prefixed literals $\circ^{k_0}L_0 \rightsquigarrow \circ^{k_1}L_1 \rightsquigarrow \dots \rightsquigarrow \circ^{k_m}L_m$ such that $k_0 = 0$, $L_0 \in \Psi$, $k_m = k$, $L_m = L'$, each k_i is between 0 and n and the \rightsquigarrow relation is defined by taking $\circ^{k_i}L_i \rightsquigarrow \circ^{k_{i+1}}L_{i+1}$ just in one of the three cases: $k_{i+1} = k_i$ and $L_i \rightarrow L_{i+1} \in \Phi$ or $k_{i+1} = k_i + 1$ and $L_i \rightarrow \circ L_{i+1} \in \Phi$ or $k_{i+1} = k_i - 1$ and $\circ L_i \rightarrow L_{i+1} \in \Phi$ (we assume that, for example, $\neg q \rightarrow \neg p \in \Phi$ whenever Φ contains $p \rightarrow q$). So, suppose conditions (**L**₁)–(**L**₂) hold. We construct an interpretation satisfying (5). By (**L**₁), both $\Psi \wedge \boxtimes \Phi$ and $\circ^n \Psi' \wedge \boxtimes \Phi$ are consistent. So, let \mathfrak{M}_{Ψ} and $\mathfrak{M}_{\Psi'}$ be such that $\mathfrak{M}_{\Psi}, 0 \models \Psi \wedge \boxtimes \Psi$ and $\mathfrak{M}_{\Psi'}, n \models \Psi' \wedge \boxtimes \Psi$, respectively. Let \mathfrak{M} be an interpretation that coincides with \mathfrak{M}_{Ψ} for all moments $k \leq 0$ and with $\mathfrak{M}_{\Psi'}$ for all $k \geq n$; for the remaining k , $0 < k < n$, it is defined as follows. First, for each $p \in \Theta$, we make p true at k and, for each $\neg p \in \Theta$, we make p false at k ; such an assignment exists due to (**L**₂). Second, we extend the assignment by making L true at k if $L \in F_{\Phi}^k(\Psi) \cup P_{\Phi}^{n-k}(\Psi')$. Observe that we have $\{p, \neg p\} \not\subseteq F_{\Phi}^k(\Psi) \cup P_{\Phi}^{n-k}(\Psi')$: for otherwise $L \wedge \boxtimes \Phi \models \circ^k p$ and $\circ^{n-k} L' \wedge \boxtimes \Phi \models \neg p$, for some $L \in \Psi$ and $L' \in \Psi'$, whence $L \wedge \boxtimes \Phi \models \circ^n \neg L'$, contrary to (**L**₁). Also, by (**L**₂), any assignment extension at this stage does not contradict the choices made due to Θ . Finally, all propositional variables not covered in the previous two cases get their values from \mathfrak{M}_{Ψ} (or $\mathfrak{M}_{\Psi'}$). We note that the last choice does not depend on the assignment that is fixed by taking account of the consequences of $\boxtimes \Phi$ with Ψ , Ψ' and Θ (because if the value of a variable depended on those sets of literals, the respective literal would be among the logical consequences and would have been fixed before). \square

Thus, it suffices to show that conditions (**L**₁) and (**L**₂) can be checked in polynomial time. First, we claim that there is a polynomial-time algorithm which, given a set Φ of binary clauses of the form $D_1 \vee D_2$, constructs a set Φ^* of binary clauses that is ‘sound and complete’ in the following sense:

- (**S**₁) $\boxtimes \Phi^* \models \boxtimes \Phi$;
 (**S**₂) if $\boxtimes \Phi \models \boxtimes (L \rightarrow \circ^k L_k)$ then either $k = 0$ and $L \rightarrow L_0 \in \Phi^*$, or $k \geq 1$ and there are L_0, L_1, \dots, L_{k-1} with $L = L_0$ and $L_i \rightarrow \circ L_{i+1} \in \Phi^*$, for $0 \leq i < k$.

Intuitively, the set Φ^* makes explicit the consequences of $\boxtimes \Phi$ and can be constructed in time $(2|\Phi|)^2$ (the number of temporal literals in Φ^* is bounded by the doubled length $|\Phi|$ of Φ as each of its literal can only be prefixed by \circ). Indeed, we start from Φ and, at each step, add $D_1 \vee D_2$ to Φ if it contains both $D_1 \vee D$ and $\neg D \vee D_2$; we also add $L_1 \vee L_2$ if Φ contains $\circ L_1 \vee \circ L_2$ (and *vice versa*). This procedure is sound since we only add consequences of $\boxtimes \Phi$; completeness follows from the completeness proof for temporal resolution [15, Section 6.3].

Our next step is to encode Φ^* by means of unary automata. Let L, L' be literals. Consider a nondeterministic finite automaton $\mathfrak{A}_{L,L'}$ over $\{0\}$ such that the literals of Φ^* are its states, with L being the initial state and L' the only accepting state, and $\{(L_1, L_2) \mid L_1 \rightarrow \circ L_2 \in \Phi^*\}$ is its transition relation. By **(S₁)** and **(S₂)**, for all $k > 0$, we have

$$\mathfrak{A}_{L,L'} \text{ accepts } 0^k \quad \text{iff} \quad \boxtimes \Phi \models \boxtimes(L \rightarrow \circ^k L').$$

Then both $F_{\Phi}^k(\Psi)$ and $P_{\Phi}^k(\Psi')$ can be defined in terms of the language of $\mathfrak{A}_{L,L'}$:

$$\begin{aligned} F_{\Phi}^k(\Psi) &= \{L' \mid \mathfrak{A}_{L,L'} \text{ accepts } 0^k, \text{ for } L \in \Psi\}, \\ P_{\Phi}^k(\Psi') &= \{L \mid \mathfrak{A}_{\neg L, \neg L'} \text{ accepts } 0^k, \text{ for } L' \in \Psi'\} \end{aligned}$$

(recall that $\circ^k L' \rightarrow L$ is equivalent to $\neg L \rightarrow \circ^k \neg L'$). Note that the numbers n and k in conditions **(L₁)** and **(L₂)** are in general exponential in the length of Φ and, therefore, the automata $\mathfrak{A}_{L,L'}$ do not immediately provide a polynomial-time procedure for checking these conditions: although it can be shown that if **(L₂)** does not hold then it fails for a polynomial number k , this is not the case for **(L₁)**, which requires the accepting state to be reached in a fixed (exponential) number of transitions. Instead, we use the *Chrobak normal form* [11] to decompose the automata into a polynomial number of polynomial-sized arithmetic progressions (which can have an exponential common period; cf. the proof of Theorem 2). In what follows, given a and b , we denote by $a + b\mathbb{N}$ the set $\{a + bm \mid m \in \mathbb{N}\}$ (the arithmetic progression with initial term a and common difference b).

It is known that every N -state unary automaton \mathfrak{A} can be converted (in polynomial time) into an equivalent automaton in Chrobak normal form (e.g., by using Martinez's algorithm [28]), which has $O(N^2)$ states and gives rise to M arithmetic progressions $a_1 + b_1\mathbb{N}, \dots, a_M + b_M\mathbb{N}$ such that

- (A₁)** $M \leq O(N^2)$ and $0 \leq a_i, b_i \leq N$, for $1 \leq i \leq M$;
- (A₂)** \mathfrak{A} accepts 0^k iff $k \in a_i + b_i\mathbb{N}$, for some $1 \leq i \leq M$.

By construction, the number of arithmetic progressions is bounded by a quadratic function in the length of Φ .

We are now in a position to give a polynomial-time algorithm for checking **(L₁)** and **(L₂)**, which requires solving Diophantine equations. In **(L₂)**, for example, to verify that, for each $p \in \Theta$, we have $\neg p \notin F_{\Phi}^k(\Psi)$, for all $0 < k < n$, we take the automata $\mathfrak{A}_{L, \neg p}$, for $L \in \Psi$, and transform them into the Chrobak normal form to obtain arithmetic progressions $a_i + b_i\mathbb{N}$, for $1 \leq i \leq M$. Then there is k , $0 < k < n$, with $\neg p \in F_{\Phi}^k(\Psi)$ iff one of the equations $a_i + b_i m = k$ has an integer solution, for some k , $0 < k < n$. The latter can be verified by taking the integer $m = \lfloor -a_i/b_i \rfloor$ and checking whether either $a_i + b_i m$ or $a_i + b_i(m + 1)$ belongs to the open interval $(0, n)$, which can clearly be done in polynomial time.

This completes the proof of Theorem 1. □

The matching lower bound for $\text{LTL}_{\text{core}}^{\square, \circ}$ -formulas, even without \square_F/\square_P , can be obtained using NP-hardness of deciding inequality of regular languages over a

1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30	
2	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0
3	1		0	1		0	1		0	1		0	1		0	1		0	1		0	1		0	1		0	1		0
5	1				0	1			0	1				0	1				0	1				0	1					0

Fig. 2. Positive numbers encoding assignments for 3 variables p_1, p_2, p_3 (shaded)

unary alphabet [27]. In the proof of Theorem 2, we give a more direct reduction of the NP-complete problem 3SAT and repeat the argument of [27, Theorem 6.1] to construct a small number of arithmetic progressions (each with a small initial term and common difference) that give rise to models of exponential size.

Theorem 2. *The satisfiability problem for $LTL_{core}^{\square, \circ}$ -formulas is NP-hard.*

Proof. The proof is by reduction of 3SAT. Let $f = \bigwedge_{i=1}^n C_i$ be a 3CNF with variables p_1, \dots, p_m and clauses C_1, \dots, C_n . By a propositional assignment for f we understand a function $\sigma: \{p_1, \dots, p_m\} \rightarrow \{0, 1\}$. We represent such assignments by sets of positive natural numbers. More precisely, let P_1, \dots, P_m be the first m prime numbers; it is known that P_m does not exceed $O(m^2)$ [1]. A natural number $k > 0$ is said to *represent* an assignment σ if k is equivalent to $\sigma(p_i)$ modulo P_i , for all $i, 1 \leq i \leq m$. Clearly, not every natural number represents an assignment since each element of

$$j + P_i \cdot \mathbb{N}, \quad \text{for } 1 \leq i \leq m \text{ and } 2 \leq j < P_i, \tag{6}$$

is equivalent to j modulo P_i with $j \geq 2$. On the other hand, every natural number that does not represent an assignment belongs to one of those arithmetic progressions (see Fig. 2).

Let C_i be a clause in f , say, $C_i = p_{i_1} \vee \neg p_{i_2} \vee p_{i_3}$. Consider

$$P_{i_1}^1 P_{i_2}^0 P_{i_3}^1 + P_{i_1} P_{i_2} P_{i_3} \cdot \mathbb{N}. \tag{7}$$

A natural number represents an assignment that makes C_i true iff it does not belong to the progressions (6) and (7). In the same way we construct a progression of the form (7) for every clause in f . Thus, a natural number $k > 0$ *does not* belong to the constructed progressions of the form (6) and (7) iff k represents a satisfying assignment for f .

To complete the proof, we show that the defined progressions can be encoded in $LTL_{core}^{\square, \circ}$. Take a propositional variable d (it will be shared by all formulas below). Given an arithmetic progression $a + b\mathbb{N}$ (with $a \geq 0$ and $b > 0$), let

$$\begin{aligned} \theta_{a,b} = & u_0 \wedge \bigwedge_{j=1}^a \boxed{\square}(u_{j-1} \rightarrow \circ_F u_j) \wedge \\ & \boxed{\square}(u_a \rightarrow v_0) \wedge \bigwedge_{j=1}^b \boxed{\square}(v_{j-1} \rightarrow \circ_F v_j) \wedge \boxed{\square}(v_b \rightarrow v_0) \wedge \boxed{\square}(v_0 \rightarrow d), \end{aligned}$$

where u_0, \dots, u_a and v_0, \dots, v_b are fresh propositional variables. It is not hard to see that, in every model of $\theta_{a,b}$, if k belongs to $a + b\mathbb{N}$, then d is true at moment k . Thus, we take a conjunction φ_f of the $\theta_{a,b}$ for arithmetic progressions (6) and (7)




Fig. 3. The minimal model of Σ_φ and \mathfrak{R}_φ

together with $p \wedge \boxtimes (\bigcirc_F p \rightarrow p) \wedge \boxtimes (p \rightarrow d) \wedge \boxtimes (\neg \boxtimes d)$, where p is a fresh variable (the last formula makes both p and d true at all moments $k \leq 0$). The size of the $\text{LTL}_{\text{core}}^{\square, \bigcirc}$ -formula φ_f is $O(n \cdot m^6)$. It is readily checked that φ_f is satisfiable iff f is satisfiable. \square

4 Core and Horn Fragments without Next-Time

Let φ be an $\text{LTL}_{\text{horn}}^{\square}$ -formula. By applying Lemma 2, we can transform φ to the form $\Psi \wedge \boxtimes \Phi^+ \wedge \boxtimes \Phi^-$, where Ψ is a set of propositional variables while Φ^+ and Φ^- are sets of *positive* and *negative clauses* of the form

$$\lambda_1 \wedge \lambda_2 \wedge \cdots \wedge \lambda_{k-1} \rightarrow \lambda_k \quad \text{and} \quad \neg \lambda_1 \vee \neg \lambda_2 \vee \cdots \vee \neg \lambda_k, \quad (8)$$

respectively. Trivially, $\Psi \wedge \boxtimes \Phi^+$ is satisfiable. Since all clauses in Φ^+ have at most one positive literal and are constructed from variables possibly prefixed by \square_F or \square_P , the formula $\Psi \wedge \boxtimes \Phi^+$ has a *canonical model* \mathfrak{R}_φ defined by taking

$$\mathfrak{R}_\varphi, n \models p \quad \text{iff} \quad \mathfrak{M}, n \models p, \quad \text{for every model } \mathfrak{M} \text{ of } \Psi \wedge \boxtimes \Phi^+, \quad n \in \mathbb{Z}$$

(indeed, $\mathfrak{R}_\varphi, 0 \models \Psi \wedge \boxtimes \Phi^+$ follows from the observation that $\mathfrak{R}_\varphi, n \models \square_F p$ iff $\mathfrak{M}, n \models \square_F p$, for every model \mathfrak{M} of $\Psi \wedge \boxtimes \Phi^+$; and similarly for $\square_P p$). If we consider the canonical model \mathfrak{R}_φ in the context of Lemma 3 then, since the language does not contain \bigcirc_F or \bigcirc_P , we have $m_{i+1} - m_i = 1$ for all i . Thus, \mathfrak{R}_φ can be thought of as a sequence of $(\ell_F - \ell_P + 1)$ -many states, the first and last of which repeat indefinitely. Let $K = |\varphi| + 4$.

Obviously, φ is satisfiable iff there is no negative clause $\neg \lambda_1 \vee \cdots \vee \neg \lambda_k$ in Φ^- such that all the λ_i are true in \mathfrak{R}_φ at some moment n with $|n| \leq K$. This condition can be encoded by means of propositional Horn clauses in the following way. For each variable p , we take $2K + 1$ variables p^n , $|n| \leq K$, and, for each $\square_F p$ and $\square_P p$, we take $2K + 1$ variables, denoted $(\square_F p)^n$ and $(\square_P p)^n$, $|n| \leq K$, respectively. Consider the following set Σ_φ of propositional Horn clauses, $|n| \leq K$:

- (H₀) p^0 , if $p \in \Psi$,
- (H₁) $\lambda_1^n \wedge \cdots \wedge \lambda_{k-1}^n \rightarrow \lambda_k^n$, if $(\lambda_1 \wedge \cdots \wedge \lambda_{k-1} \rightarrow \lambda_k) \in \Phi^+$,
- (H₂) $(\square_F p)^n \rightarrow (\square_F p)^{n+1}$ if $n < K$, $(\square_F p)^n \rightarrow (\square_F p)^{n-1}$ if $n > -K$,
- (H₃) $(\square_P p)^n \rightarrow p^{n+1}$, $(\square_P p)^n \rightarrow p^{n-1}$,
- (H₄) $(\square_F p)^n \wedge p^n \rightarrow (\square_F p)^{n-1}$ if $n > -K$, $(\square_F p)^n \wedge p^n \rightarrow (\square_F p)^{n+1}$ if $n < K$,
- (H₅) $(\square_F p)^K \leftrightarrow p^K$, $(\square_P p)^{-K} \leftrightarrow p^{-K}$,
- (H₆) $(\square_F p)^{-K} \leftrightarrow p^{-K}$, $(\square_P p)^K \leftrightarrow p^K$.

Clearly, $|\Sigma_\varphi| \leq O(|\varphi|^2)$. It is readily seen that the minimal model of Σ_φ corresponds to the canonical model \mathfrak{K}_φ as shown in Fig. 3. As propositional Horn satisfiability is PTIME-complete, we obtain the following:

Theorem 3. *The satisfiability problem for $\text{LTL}_{\text{horn}}^\square$ -formulas is in PTIME.*

4.1 Temporal Derivations for $\text{LTL}_{\text{core}}^\square$ in NLogSpace

In $\text{LTL}_{\text{core}}^\square$ -formulas, all clauses are binary: $k = 2$ in (8). Satisfiability of propositional binary clauses is known to be NLOGSPACE-complete. However, in the reduction $\varphi \mapsto \Sigma_\varphi$ above, the clauses (\mathbf{H}_4) are ternary. In this section we show how to modify the reduction to ensure membership in NLOGSPACE. More precisely, we define two types of derivation from $\Psi \wedge \boxtimes \Phi^+$: a 0-derivation of (λ, n) will mean that $\mathfrak{K}_\varphi, n \models \lambda$, while a \forall -derivation of λ from λ' that $\mathfrak{K}_\varphi, 0 \models \boxtimes \lambda' \rightarrow \boxtimes \lambda$. We then show that these derivations define \mathfrak{K}_φ and that satisfiability of φ can be checked by a nondeterministic algorithm in logarithmic space.

Denote by \rightarrow^* the transitive and reflexive closure of the relation \rightarrow over literals given by the clauses of Φ^+ . We require the following derivation rules over the pairs (λ, n) , where λ is a positive temporal literal in φ and $n \in \mathbb{Z}$:

- (**R**₁) $(\lambda_1, n) \Rightarrow (\lambda_2, n), \quad \text{if } \lambda_1 \rightarrow^* \lambda_2,$
- (**R**₂) $(\Box_F p, n) \Rightarrow (\Box_F p, n + 1), \quad (\Box_F p, n) \Rightarrow (\Box_F p, n - 1),$
- (**R**₃) $(\Box_F p, n) \Rightarrow (p, n + 1), \quad (\Box_F p, n) \Rightarrow (p, n - 1),$
- (**R**₄) $(\Box_F p, 0) \Rightarrow (\Box_F p, -1), \quad (\Box_F p, 0) \Rightarrow (\Box_F p, 1), \quad \text{if } p' \rightarrow^* p \text{ for } p' \in \Psi,$
- (**R**₅) $(p, n) \Rightarrow (\Box_F p, n - 1), \quad (p, n) \Rightarrow (\Box_F p, n + 1).$

The rules in (**R**₁)–(**R**₄) mimic (**H**₁)–(**H**₄) above ((**H**₄) at moment 0 only) and reflect the semantics of LTL in the sense that whenever $(\lambda, n) \Rightarrow (\lambda', n')$ and $\mathfrak{K}_\varphi, n \models \lambda$ then $\mathfrak{K}_\varphi, n' \models \lambda'$. For example, consider (**R**₄). It only applies if p follows (by \rightarrow^*) from the initial conditions in Ψ , in which case $\mathfrak{K}_\varphi, 0 \models p$, and so $\mathfrak{K}_\varphi, 0 \models \Box_F p$ implies $\mathfrak{K}_\varphi, -1 \models \Box_F p$. The rules in (**R**₅) are different: for instance, we can only apply $(p, n) \Rightarrow (\Box_F p, n - 1)$ if we know that p holds at all $m \geq n$.

A sequence $\mathfrak{d}: (\lambda_0, n_0) \Rightarrow \dots \Rightarrow (\lambda_\ell, n_\ell)$, for $\ell \geq 0$, is called a 0-derivation of (λ_ℓ, n_ℓ) if $\lambda_0 \in \Psi$, $n_0 = 0$ and all applications of (**R**₅) are safe in the following sense: for any $(p, n_i) \Rightarrow_{(\mathbf{R}_5)} (\Box_F p, n_i - 1)$, there is $\lambda_j = \Box_F q$, for some q and $0 \leq j < i$; similarly, for any $(p, n_i) \Rightarrow_{(\mathbf{R}_5)} (\Box_F p, n_i + 1)$, there is $\lambda_j = \Box_F q$ with $0 \leq j < i$. In this case we write $\Psi \Rightarrow^0 (\lambda_\ell, n_\ell)$. For example, consider

$$\varphi = p \wedge \boxtimes(p \rightarrow \Box_F q) \wedge \boxtimes(q \rightarrow r) \wedge \boxtimes(p \rightarrow r).$$

Evidently, $\mathfrak{K}_\varphi, -1 \models \Box_F r$. The following sequence is a 0-derivation of $(\Box_F r, -1)$ because the application of (**R**₅) is safe due to $\Box_F q$:

$$(p, 0) \Rightarrow_{(\mathbf{R}_1)} (\Box_F q, 0) \Rightarrow_{(\mathbf{R}_3)} (q, 1) \Rightarrow_{(\mathbf{R}_1)} (r, 1) \Rightarrow_{(\mathbf{R}_5)} (\Box_F r, 0) \Rightarrow_{(\mathbf{R}_4)} (\Box_F r, -1).$$

Intuitively, if we can derive $(r, 1)$ using $(\Box_F q, 0)$, then we can also derive (r, n) for any $n \geq 1$, and so we must also have $(\Box_F r, 0)$, which justifies the application of (**R**₅). This argument is formalised in the following lemma:




Fig. 4. Removing applications of (\mathbf{R}_4) (left) and shifting a 0-derivation by 2 (right): dashed arrows show the original derivation and solid ones the resulting derivation

Lemma 5 (monotonicity). *Let \mathfrak{d} be a 0-derivation of (λ_ℓ, n_ℓ) with a suffix*

$$\mathfrak{s}: (\Box_F q, n_s) \Rightarrow (\lambda_{s+1}, n_{s+1}) \Rightarrow \cdots \Rightarrow (\lambda_\ell, n_\ell), \quad (9)$$

where none of the λ_i contains \Box_F . Then $\Psi \Rightarrow^0 (\lambda_\ell, m)$, for all $m \geq n_\ell$. Similarly, if there is a suffix beginning with some $\Box_P q$ then $\Psi \Rightarrow^0 (\lambda_\ell, m)$, for all $m \leq n_\ell$. Moreover, these 0-derivations only contain the rules used in \mathfrak{d} and (\mathbf{R}_2) .

Proof. We first remove all applications of (\mathbf{R}_4) in \mathfrak{s} . Let $(\lambda_i, n_i) \Rightarrow_{(\mathbf{R}_4)} (\lambda_{i+1}, n_{i+1})$ be the first one. By definition, $n_i = 0$ and, since $\Box_F q$ is the last \Box_F in \mathfrak{d} , we have $n_{i+1} = 1$ and $\lambda_i = \lambda_{i+1} = \Box_P r$, for some r . So we can begin \mathfrak{s} with $(\Box_F q, n_s) \Rightarrow_{(\mathbf{R}_2)} (\Box_F q, n_s + 1) \Rightarrow (\lambda_{s+1}, n_{s+1} + 1) \Rightarrow \cdots \Rightarrow (\lambda_i, n_i + 1) \Rightarrow (\lambda_{i+2}, n_{i+2})$; see Fig. 4 on the left-hand side. We repeatedly apply this operation to obtain a suffix \mathfrak{s} of the form (9) that does not use (\mathbf{R}_4) . We then replace \mathfrak{s} in \mathfrak{d} with $(\Box_F q, n_s) \Rightarrow_{(\mathbf{R}_2)} \cdots \Rightarrow_{(\mathbf{R}_2)} (\Box_F q, n_s + k) \Rightarrow (\lambda_{s+1}, n_{s+1} + k) \Rightarrow \cdots \Rightarrow (\lambda_\ell, n_\ell + k)$, where $k = m - n_\ell$; see Fig. 4 on the right-hand side. \square

However, 0-derivations are not enough to obtain all literals that are true in \mathfrak{K}_φ . Indeed, consider the formula

$$\varphi = r \wedge \boxtimes(r \rightarrow \Box_F q) \wedge \boxtimes(\Box_F q \rightarrow q) \wedge \boxtimes(\Box_F q \rightarrow p).$$

Clearly, $\mathfrak{K}_\varphi, n \models p$ for all $n \in \mathbb{Z}$, but neither (p, n) nor $(\Box_P q, n)$ is 0-derivable. On the other hand, for each $n \in \mathbb{Z}$, there is a 0-derivation of (q, n) : for example,

$$(r, 0) \Rightarrow_{(\mathbf{R}_1)} (\Box_F q, 0) \Rightarrow_{(\mathbf{R}_1)} (q, 0) \Rightarrow_{(\mathbf{R}_5)} (\Box_F q, -1) \Rightarrow_{(\mathbf{R}_1)} (q, -1).$$

These 0-derivations correspond to $\mathfrak{K}_\varphi, 0 \models \boxtimes q$, from which we can derive $\boxtimes p$ by means of the second type of derivations. A sequence $\mathfrak{d}: (\lambda_0, n_0) \Rightarrow \cdots \Rightarrow (\lambda_\ell, n_\ell)$ is called a \forall -derivation of λ_ℓ from λ_0 if it uses only (\mathbf{R}_1) – (\mathbf{R}_3) and (\mathbf{R}_5) , whose applications are not necessarily safe. So we write $\Psi \Rightarrow^\forall \lambda$ if there is a \forall -derivation of λ from some q such that $\Psi \Rightarrow^0 (q, n)$, for all $n \in \mathbb{Z}$. In the example above, $(q, 0) \Rightarrow_{(\mathbf{R}_5)} (\Box_P q, 1) \Rightarrow_{(\mathbf{R}_1)} (p, 1)$ is a \forall -derivation of p from q , whence $\Psi \Rightarrow^\forall p$.

Lemma 6 (soundness). *If $\Psi \Rightarrow^0 (\lambda, n)$ then $\mathfrak{R}_\varphi, n \models \lambda$. If $\Psi \Rightarrow^\forall \lambda$ then $\mathfrak{R}_\varphi, 0 \models \boxplus \lambda$.*

Proof. By induction on the derivation length, using Lemma 5 for **(R₅)**. \square

Lemma 7 (completeness). *If $\mathfrak{R}_\varphi, n \models \lambda$ then either $\Psi \Rightarrow^0 (\lambda, n)$ or $\Psi \Rightarrow^\forall \lambda$.*

Proof. Let \mathfrak{M} be an interpretation such that, for all p and $n \in \mathbb{Z}$, we have $\mathfrak{M}, n \models p$ iff $\Psi \Rightarrow^0 (p, n)$ or $\Psi \Rightarrow^\forall p$. It suffices to show that $\mathfrak{M}, 0 \models \Psi \wedge \boxplus \Phi^+$. Indeed, if we assume that there are p' and n' such that $\mathfrak{R}_\varphi, n' \models p'$ but neither $\Psi \Rightarrow^0 (p', n')$ nor $\Psi \Rightarrow^\forall p'$, we will obtain $\mathfrak{M}, n' \models \neg p'$ contrary to our assumption (other types of literals are considered analogously).

Thus, we have to show that \mathfrak{M} is a model of $\Psi \wedge \boxplus \Phi^+$. Suppose $p \in \Psi$. Then trivially $\Psi \Rightarrow^0 (p, 0)$, and so $\mathfrak{M}, 0 \models p$. Suppose $\lambda_1 \rightarrow \lambda_2 \in \Phi^+$ and $\mathfrak{M}, n \models \lambda_1$. We consider three cases depending on the shape of λ_1 and show that $\mathfrak{M}, n \models \lambda_2$.

$\lambda_1 = p$. If $\Psi \Rightarrow^\forall p$ then, by **(R₁)**, $\Psi \Rightarrow^\forall \lambda_2$. Otherwise, there is a 0-derivation of (p, n) , and so $\Psi \Rightarrow^0 (\lambda_1, n) \Rightarrow_{\mathbf{(R_1)}} (\lambda_2, n)$.

$\lambda_1 = \square_F p$. Then $\mathfrak{M}, m \models p$ for all $m > n$. Consider $\mathfrak{M}, n+1 \models p$. If $\Psi \Rightarrow^\forall p$ then, by **(R₅)**, **(R₁)**, $\Psi \Rightarrow^\forall \lambda_2$. Otherwise, there is a 0-derivation \mathfrak{d} of $(p, n+1)$.

(F) If \square_F occurs in \mathfrak{d} then $\Psi \Rightarrow^0 (p, n+1) \Rightarrow_{\mathbf{(R_5)}} (\square_F p, n) \Rightarrow_{\mathbf{(R_1)}} (\lambda_2, n)$.

(P) If \square_P occurs in \mathfrak{d} then, by Lemma 5, $\Psi \Rightarrow^0 (p, m)$ for each $m \leq n+1$.

Thus, $\Psi \Rightarrow^0 (p, m)$ for all $m \in \mathbb{Z}$, and so, by **(R₅)** and **(R₁)**, $\Psi \Rightarrow^\forall \lambda_2$.

(0) If \mathfrak{d} contains neither \square_F nor \square_P then $n = -1$ and $\lambda \rightarrow^* p$, for some $\lambda \in \Psi$ (by **(R₁)**). As $\mathfrak{M}, 1 \models p$ and we assumed $\Psi \not\Rightarrow^\forall p$, there is a 0-derivation \mathfrak{d}' of $(p, 1)$, which must contain \square_F or \square_P . If \mathfrak{d}' contains \square_F then $\Psi \Rightarrow^0 (p, 1) \Rightarrow_{\mathbf{(R_5)}} (\square_F p, 0) \Rightarrow_{\mathbf{(R_4)}} (\square_F p, -1) \Rightarrow_{\mathbf{(R_1)}} (\lambda_2, n)$. If \square_P occurs in \mathfrak{d}' then, by the argument in (P), $\Psi \Rightarrow^\forall \lambda_2$.

$\lambda_1 = \square_P p$. The proof is symmetric.

In each of these cases, we have either $\Psi \Rightarrow^0 (\lambda_2, n)$ or $\Psi \Rightarrow^\forall \lambda_2$. Observe that $\Psi \Rightarrow^0 (\lambda_2, n)$ implies $\mathfrak{M}, n \models \lambda_2$. Indeed, this clearly holds for $\lambda_2 = p$. If $\lambda_2 = \square_F p$ then, by repetitive applications of **(R₂)** and an application of **(R₃)**, we obtain $\Psi \Rightarrow^0 (p, m)$, for all $m > n$, which means $\mathfrak{M}, n \models \square_F p$. The case $\lambda_2 = \square_P p$ is symmetric. If $\Psi \Rightarrow^\forall \lambda_2$ then, independently of whether λ_2 is p' , $\square_F p'$ or $\square_P p'$, we have $\Psi \Rightarrow^\forall p'$, so $\mathfrak{M}, m \models p'$ for all $m \in \mathbb{Z}$, whence, $\mathfrak{M}, n \models \lambda_2$. \square

Next, in Lemmas 8 and 9, we provide efficient criteria for checking the conditions $\Psi \Rightarrow^0 (\lambda, n)$ and $\Psi \Rightarrow^\forall \lambda$ by restricting the range of numbers that can be used in 0-derivations (numbers in \forall -derivations can simply be ignored). Given a 0-derivation $\mathfrak{d}: (\lambda_0, n_0) \Rightarrow \dots \Rightarrow (\lambda_\ell, n_\ell)$, we define its *reach* as

$$r(\mathfrak{d}) = \max\{|n_i| \mid 0 \leq i \leq \ell\}.$$

We say that \mathfrak{d} *right-stutters*, if there are $v < w$ such that $\lambda_v = \lambda_w$, $n_v < n_w$ and $n_i > 0$, for all i , $v \leq i \leq w$ (in particular, **(R₄)** is not applied between v and w). Symmetrically, \mathfrak{d} *left-stutters* if there are $v < w$ such that $\lambda_v = \lambda_w$, $n_v > n_w$ and $n_i < 0$, for all i , $v \leq i \leq w$.




Fig. 5. Left-stuttering: n_v and n_w occur between -1 and $-|\varphi| - 1$ (shaded) and the fragment of the derivation from n_v to n_w can be repeated any number of times (including 0)

Lemma 8 (checking \Rightarrow^0). $\Psi \Rightarrow^0 (\lambda, n)$ iff there exists a 0-derivation \mathfrak{d} of (λ, m) such that $r(\mathfrak{d}) \leq 2|\varphi|$ and one of the following conditions holds:

- (C₁) $m = n$;
- (C₂) \mathfrak{d} contains \square_F and either $m \leq n$ or \mathfrak{d} left-stutters;
- (C₃) \mathfrak{d} contains \square_P and either $m \geq n$ or \mathfrak{d} right-stutters.

Proof. (\Rightarrow) Let $\mathfrak{d}: (\lambda_0, n_0) \Rightarrow \dots \Rightarrow (\lambda_\ell, n_\ell)$ be a 0-derivation of (λ, n) . If $r(\mathfrak{d}) \leq |\varphi|$ then \mathfrak{d} satisfies (C₁). Otherwise, we take the first \square -literal in \mathfrak{d} , say $\lambda_t = \square_F q$ (the case of $\square_P q$ is symmetric). Clearly, $|n_t| \leq 1$. Let $u > t$ be the smallest index with $|n_u| > |\varphi|$. Since adjacent n_i and n_{i+1} differ by at most 1, the segment between (λ_t, n_t) and (λ_u, n_u) contains a repeating literal: more precisely, there exist $v < w$ between t and u such that $\lambda_v = \lambda_w$ and

- either $n_v > n_w$ and $n_i < 0$, for $v \leq i \leq w$,
- or $n_v < n_w$ and $n_i > 0$, for $v \leq i \leq w$.

In the former case \mathfrak{d} left-stutters, and we perform the following operations on the suffix $\mathfrak{s}: (\lambda_w, n_w) \Rightarrow \dots \Rightarrow (\lambda_\ell, n_\ell)$ of \mathfrak{d} . First, we eliminate all applications of (R₄) in \mathfrak{s} : each suffix $(\square_F q, 0) \Rightarrow_{(R_4)} (\square_F q, -1) \Rightarrow (\lambda_s, n_s) \Rightarrow \dots \Rightarrow (\lambda_\ell, n_\ell)$ is replaced by $(\square_F q, 0) \Rightarrow (\lambda_s, n_s + 1) \Rightarrow \dots \Rightarrow (\lambda_\ell, n_\ell + 1)$; and similarly for \square_P . If each time we eliminate the last application of (R₄) then the result is clearly a 0-derivation. Second, we remove all duplicating literals: each suffix $(\lambda_s, n_s) \Rightarrow \dots \Rightarrow (\lambda_{s'}, n_{s'}) \Rightarrow (\lambda_{s'+1}, n_{s'+1}) \Rightarrow \dots \Rightarrow (\lambda_\ell, n_\ell)$ with $\lambda_s = \lambda_{s'}$ is replaced by $(\lambda_s, n_s) \Rightarrow (\lambda_{s'+1}, n_{s'+1} + k) \Rightarrow \dots \Rightarrow (\lambda_\ell, n_\ell + k)$, where $k = n_s - n_{s'}$. This will give us a left-stuttering 0-derivation \mathfrak{d}' of (λ, m) , for some m . Since there are at most $|\varphi|$ distinct literals in \mathfrak{s} , we have $r(\mathfrak{d}') \leq 2|\varphi|$, thus satisfying the second option of (C₂); see Fig. 5.

In the latter case \mathfrak{d} right-stutters, and we construct a 0-derivation \mathfrak{d}' of (p, n') by cutting out the segment $(\lambda_{v+1}, n_{v+1}) \Rightarrow \dots \Rightarrow (\lambda_w, n_w)$ from \mathfrak{d} and ‘shifting’ the tail using the construction above: eliminate applications of (R₄) and then decrease all numbers by $n_w - n_v > 0$. We then consider the obtained \mathfrak{d}' as the original \mathfrak{d} . As the length of the derivations decreases and $n' \leq n$, by applying

this procedure sufficiently many times, we shall finally construct a 0-derivation of reach $\leq 2|\varphi|$ and satisfying either (C_1) or the first option of (C_2) .

(\Leftarrow) is left to the reader. \square

In a similar way we can show how to efficiently check the condition $\Psi \Rightarrow^\forall p$:

Lemma 9 (checking \Rightarrow^\forall). $\Psi \Rightarrow^0 (\lambda, n)$ holds for all $n \in \mathbb{Z}$ iff there are 0-derivations \mathfrak{d} of (λ, m) and \mathfrak{d}' of (λ, m') of reach at most $2|\varphi|$ such that one of the following conditions holds:

(C'_1) \mathfrak{d} contains \square_F , \mathfrak{d}' contains \square_P and $m \leq m' + 1$;

(C'_2) \mathfrak{d} contains \square_F and left-stutters;

(C'_3) \mathfrak{d} contains \square_P and right-stutters.

Proof. (\Rightarrow) Take a 0-derivation of $(q, 2|\varphi| + 1)$. By Lemma 8, there is a derivation \mathfrak{d}_0 of (q, n_0) with $r(\mathfrak{d}_0) \leq 2|\varphi|$ satisfying either (C_2) or (C_3) . If \mathfrak{d}_0 left- or right-stutters then we have (C'_2) or (C'_3) , respectively. Otherwise, \mathfrak{d}_0 contains \square_F and we can construct a finite sequence of 0-derivations $\mathfrak{d}_0, \mathfrak{d}_1, \mathfrak{d}_2, \dots, \mathfrak{d}_k$ of reach at most $2|\varphi|$, where each \mathfrak{d}_i is a 0-derivation of (q, n_i) containing \square_F , and such that $n_0 > n_1 > n_2 > \dots > n_k$.

Suppose we have already constructed \mathfrak{d}_i . Since $\Psi \Rightarrow^0 (q, n)$, for all n , we have $\Psi \Rightarrow^0 (q, n_i - 1)$. By Lemma 8, there is a 0-derivation \mathfrak{d} of (q, n_{i+1}) , for some n_{i+1} , with one of (C_1) – (C_3) . If (C_2) and \mathfrak{d} left-stutters or (C_3) and \mathfrak{d} right-stutters then we obtain (C'_2) or (C'_3) , respectively. If (C_2) and \mathfrak{d} contains \square_F with $n_{i+1} \leq n_i - 1$ then \mathfrak{d} becomes the next member \mathfrak{d}_{i+1} in the sequence. If (C_3) and \mathfrak{d} contains \square_P with $n_{i+1} \geq n_i - 1$ then \mathfrak{d}_i and \mathfrak{d} satisfy (C'_1) . Otherwise, we have (C_1) with $n_{i+1} = n_i - 1$ (recall that $n_i > -2|\varphi|$). Consider three cases. If \mathfrak{d} contains \square_F then \mathfrak{d} becomes the next member \mathfrak{d}_{i+1} in the sequence. If \mathfrak{d} contains \square_P then \mathfrak{d}_i and \mathfrak{d} satisfy (C'_1) . Otherwise, that is, if \mathfrak{d} contains neither \square_P nor \square_F , we must have $n_{i+1} = 0$ and $p \rightarrow^* q$, for some $p \in \Psi$. Then we have $n_i = 1$ and, as \mathfrak{d}_i contains \square_F , we can append $(q, 1) \Rightarrow_{(R_5)} (\square_F q, 0) \Rightarrow_{(R_4)} (\square_F q, -1) \Rightarrow_{(R_3)} (q, 0)$ to \mathfrak{d} to obtain the next member \mathfrak{d}_{i+1} in the sequence.

(\Leftarrow) is left to the reader. \square

We are now in a position to prove the main result of this section.

Theorem 4. *The satisfiability problem for LTL_{core}^\square -formulas is in NLOGSPACE.*

Proof. An LTL_{core}^\square -formula $\varphi = \Psi \wedge \boxtimes \Phi^+ \wedge \boxtimes \Phi^-$ is unsatisfiable iff Φ^- contains a clause $\neg\lambda_1 \vee \neg\lambda_2$ such that $\mathfrak{R}_\varphi, n \models \lambda_1 \wedge \lambda_2$, for some n with $|n| \leq K$. For each $\neg\lambda_1 \vee \neg\lambda_2$ in Φ^- , our algorithm guesses such an n (in binary) and, for both λ_1 and λ_2 , checks whether $\Psi \Rightarrow^0 (\lambda_i, n)$ or $\Psi \Rightarrow^\forall \lambda_i$, which, by Lemmas 8 and 9, requires only logarithmic space. \square

The initial clauses of LTL_{core}^\square -formulas φ are propositional variables. If we slightly extend the language to allow for initial core-clauses (without \boxtimes), then the satisfiability problem becomes PTIME-hard. This can be shown by reduction of satisfiability of propositional Horn formulas with clauses of the form $p, \neg p$ and

$p \wedge q \rightarrow r$, which is known to be PTIME-complete. Indeed, suppose $f = \bigwedge_{i=1}^n C_i$ is such a formula. We define a temporal formula φ_f to be the conjunction of all unary clauses of f with the following formulas, for each ternary clause C_i of the form $p \wedge q \rightarrow r$ in f :

$$c_i \wedge \boxed{(p \rightarrow \square_F c_i)} \wedge \boxed{(q \rightarrow \square_F c_i)} \wedge (\boxed{c_i} \rightarrow r),$$

where c_i is a fresh variable. One can show that f is satisfiable iff φ_f is satisfiable.

We finish this section by an observation that if the language allows for non-Horn clauses (e.g., $p \vee q$) then the satisfiability problem becomes NP-hard:

Theorem 5. *The satisfiability problem for $\text{LTL}_{krom}^{\square}$ -formulas is NP-hard.*

Proof. By reduction of graph 3-colourability. Given a graph $G = (V, E)$, consider the following $\text{LTL}_{krom}^{\square}$ -formula φ_G with variables p_0, \dots, p_4 and \bar{v}_i , for $v_i \in V$:

$$p_0 \wedge \bigwedge_{0 \leq i \leq 3} \boxed{(p_i \rightarrow \square_F p_{i+1})} \wedge \bigwedge_{v_i \in V} \boxed{(p_0 \rightarrow \neg \square_F \bar{v}_i)} \wedge \\ \bigwedge_{v_i \in V} \boxed{(p_4 \rightarrow \bar{v}_i)} \wedge \bigwedge_{(v_i, v_j) \in E} \boxed{(\bar{v}_i \vee \bar{v}_j)}.$$

Intuitively, the first four conjuncts of this formula choose, for each vertex v_i of the graph, a moment of time $1 \leq n_i \leq 3$; the last conjunct makes sure that $n_i \neq n_j$ in case v_i and v_j are connected by an edge in G . It can be easily shown that φ_G is satisfiable iff G is 3-colourable. \square

5 Conclusion

We have investigated the computational complexity of the satisfiability problem for the fragments of LTL over $(\mathbb{Z}, <)$ given by the form of the clauses—*bool*, *horn*, *krom* and *core*—in the clausal normal form and the temporal operators available for constructing temporal literals. Apart from $\text{LTL}_{bool}^{\square, \circ}$, whose formulas are equisatisfiable to formulas in the full LTL, only $\text{LTL}_{horn}^{\square, \circ}$ has PSPACE-complete satisfiability. For all other fragments, the complexity varies from NLOGSPACE to PTIME and NP.

The idea to consider sub-Boolean fragments of LTL comes from description logic, where the *DL-Lite* family [9,3] of logics has been designed and investigated with the aim of finding formalisms suitable for ontology-based data access (OBDA). It transpired that, despite their low complexity, *DL-Lite* logics were capable of representing basic conceptual data modelling constructs [8,2], and gave rise to the W3C standard ontology language *OWL 2 QL* for OBDA. One possible application of the results obtained in this paper lies in temporal conceptual modelling and temporal OBDA [5]. Temporal description logics (and other many-dimensional logics) are notorious for their bad computational properties [17,20]. We believe, however, that efficient practical reasoning can be achieved by considering sub-Boolean temporal extensions of *DL-Lite* logics; see [4] for first promising results.

References

1. Apostol, T.: *Introduction to Analytic Number Theory*. Springer (1976)
2. Artale, A., Calvanese, D., Kontchakov, R., Ryzhikov, V., Zakharyashev, M.: Reasoning over extended ER models. In: Parent, C., Schewe, K.-D., Storey, V.C., Thalheim, B. (eds.) *ER 2007*. LNCS, vol. 4801, pp. 277–292. Springer, Heidelberg (2007)
3. Artale, A., Calvanese, D., Kontchakov, R., Zakharyashev, M.: The DL-Lite family and relations. *Journal of Artificial Intelligence Research* 36, 1–69 (2009)
4. Artale, A., Kontchakov, R., Ryzhikov, V., Zakharyashev, M.: Past and future of DL-Lite. In: *Proc. of AAAI*, pp. 243–248 (2010)
5. Artale, A., Kontchakov, R., Ryzhikov, V., Zakharyashev, M.: Complexity of reasoning over temporal data models. In: Parsons, J., Saeki, M., Shoval, P., Woo, C., Wand, Y. (eds.) *ER 2010*. LNCS, vol. 6412, pp. 174–187. Springer, Heidelberg (2010)
6. Artale, A., Kontchakov, R., Ryzhikov, V., Zakharyashev, M.: The Complexity of Clausal Fragments of LTL. *CoRR* abs/1306.5088 (2013)
7. Bauland, M., Schneider, T., Schnoor, H., Schnoor, I., Vollmer, H.: The complexity of generalized satisfiability for linear temporal logic. *Logical Methods in Computer Science* 5(1) (2009)
8. Berardi, D., Calvanese, D., De Giacomo, G.: Reasoning on UML class diagrams. *Artificial Intelligence* 168(1-2), 70–118 (2005)
9. Calvanese, D., De Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Tractable reasoning and efficient query answering in description logics: The DL-Lite family. *Journal of Automated Reasoning* 39(3), 385–429 (2007)
10. Chen, C.-C., Lin, I.-P.: The computational complexity of satisfiability of temporal Horn formulas in propositional linear-time temporal logic. *Information Processing Letters* 45(3), 131–136 (1993)
11. Chrobak, M.: Finite automata and unary languages. *Theoretical Computer Science* 47(2), 149–158 (1986)
12. Demri, S., Schnoebelen, P.: The complexity of propositional linear temporal logics in simple cases. *Information and Computation* 174(1), 84–103 (2002)
13. Dixon, C., Fisher, M., Konev, B.: Tractable temporal reasoning. In: *Proc. of IJCAI*, pp. 318–323 (2007)
14. Fisher, M.: A resolution method for temporal logic. In: *Proc. of IJCAI*, pp. 99–104. Morgan Kaufmann (1991)
15. Fisher, M., Dixon, C., Peim, M.: Clausal temporal resolution. *ACM Transactions on Computational Logic* 2(1), 12–56 (2001)
16. Gabbay, D., Hodkinson, I., Reynolds, M.: *Temporal Logic: Mathematical Foundations and Computational Aspects*, vol. 1. Oxford University Press (1994)
17. Gabbay, D., Kurucz, A., Wolter, F., Zakharyashev, M.: *Many-Dimensional Modal Logics: Theory and Applications*. Elsevier (2003)
18. Halpern, J., Reif, J.: The propositional dynamic logic of deterministic, well-structured programs. In: *Proc. of FOCS*, pp. 322–334. IEEE (1981)
19. Lichtenstein, O., Pnueli, A., Zuck, L.D.: The glory of the past. In: Parikh, R. (ed.) *Logic of Programs 1985*. LNCS, vol. 193, pp. 196–218. Springer, Heidelberg (1985)
20. Lutz, C., Wolter, F., Zakharyashev, M.: Temporal description logics: A survey. In: *Proc. of TIME*, pp. 3–14. IEEE Comp. Society (2008)
21. Markey, N.: Past is for free: On the complexity of verifying linear temporal properties with past. *Acta Informatica* 40(6-7), 431–458 (2004)

22. Ono, H., Nakamura, A.: On the size of refutation Kripke models for some linear modal and tense logics. *Studia Logica* 39, 325–333 (1980)
23. Plaisted, D.: A decision procedure for combinations of propositional temporal logic and other specialized theories. *Journal of Automated Reasoning* 2, 171–190 (1986)
24. Rabinovich, A.: Temporal logics over linear time domains are in PSPACE. In: Kučera, A., Potapov, I. (eds.) RP 2010. LNCS, vol. 6227, pp. 29–50. Springer, Heidelberg (2010)
25. Reynolds, M.: The complexity of decision problems for linear temporal logics. *Journal of Studies in Logic* 3(1), 19–50 (2010)
26. Sistla, A., Clarke, E.: The complexity of propositional linear temporal logics. In: Proc. of STOC, pp. 159–168. ACM (1982)
27. Stockmeyer, L., Meyer, A.: Word problems requiring exponential time: Preliminary report. In: Proc. of STOC, pp. 1–9. ACM (1973)
28. To, A.W.: Unary finite automata vs. arithmetic progressions. *Information Processing Letters* 109(17), 1010–1014 (2009)