

Computing All Facts Entailed By An LTL Specification

Przemysław Andrzej Wałęga¹, Michał Zawidzki^{1,2}, Christoph Haase¹

¹Department of Computer Science, University of Oxford

²Department of Logic, University of Łódź

{przemyslaw.walega, michal.zawidzki, christoph.haase}@cs.ox.ac.uk

Abstract

We study the problem of efficiently computing all (usually infinitely many) facts which are entailed by a specification written in linear temporal logic (LTL)—a standard formalism for specifying and verifying properties of computations in reactive systems. This problem can be seen as a generalisation of the standard entailment checking, but whose output provides a much wider understanding of the system’s behaviour. We show that in full LTL the problem can be solved in doubly exponential time, whereas for Horn fragments of LTL, which can be seen as temporal logic programs, the problem can be solved in exponential or only quadratic time, depending on the allowed temporal operators in the input formula. Moreover, we show that these bounds are optimal. We also implement and experimentally compare two techniques for solving the problem: an automata-based algorithm for full LTL and a materialisation-based algorithm for Horn fragments. The obtained results suggest practical usefulness of our approach.

1 Introduction

Analysing temporal data is of increasing importance due to a number of emerging applications such as modelling blockchain-based smart contracts (Alharby and Van Moorsel 2017), identifying traffic anomalies (Münz and Carle 2007), algorithmic trading (Nutti et al. 2011), and IoT applications (Lécué 2017). To capture formally the evolving behaviour of dynamic domains, various temporal languages have been introduced (Vardi 2009; Abadi and Manna 1989; Demri, Goranko, and Lange 2016), including the seminal linear temporal logic (LTL) (Pnueli 1977) with its fragments and modifications, see, e.g., Alur and La Torre (2004), Artale et al. (2013), Demri and Schnoebelen (2002).

LTL-based formalisms are currently commonly used in KRR to express time-sensitive knowledge and ontologies (Artale et al. 2021; Lutz, Wolter, and Zakharyashev 2008; Abadi and Manna 1989; Aguado et al. 2013; Ronca et al. 2018), as illustrated in the following simple example.

Example 1. Consider a system consisting of 3 devices with $IDs = \{1, 2, 3\}$, all initially backed up on the same day, but each one in some unknown mode from the set $M = \{1, \dots, 10\}$. We express it with Formulas (1) for all $id \in IDs$, where proposition $Bckp_m^{id}$ being true at some day represents the fact that the device number id is backed up on

this day in mode m . We assume that no device is simultaneously backed up in two different modes m and m' , as stated by Formulas (2) for all $id \in IDs$ and distinct $m, m' \in M$, where G is the ‘always in the future’ LTL-operator. Modes determine the frequency of backups, namely backups in a mode $m \in M$ are performed regularly every m days, as stated by Formulas (3) with m iterations of the ‘next’ operator X . The system is safe on a given day if all three devices are backed up on this day, as stated by Formulas (4), for all $m, m', m'' \in M$.

$$Bckp_1^{id} \vee \dots \vee Bckp_{10}^{id}, \quad (1)$$

$$G (Bckp_m^{id} \wedge Bckp_{m'}^{id} \rightarrow \perp), \quad (2)$$

$$G (Bckp_m^{id} \rightarrow X^m Bckp_m^{id}), \quad (3)$$

$$G (Bckp_m^1 \wedge Bckp_{m'}^2 \wedge Bckp_{m''}^3 \rightarrow Safe). \quad (4)$$

The main reasoning problems considered in this setting are checking consistency (satisfiability) of the input temporal knowledge and answering temporal queries mediated by a temporal knowledge (formula entailment). There exist various approaches to solve these problems, including tableau systems (Wolper 1985; Lichtenstein and Pnueli 2000), automata-based techniques (Vardi 2005; Li et al. 2014; Duret-Lutz et al. 2022), reductions to SAT (Geatti, Gigante, and Montanari 2021), and rewriting to first-order logic (Artale et al. 2021; Ryzhikov, Wałęga, and Zakharyashev 2020), which have been implemented in a plethora of reasoning systems. Applying these methods, for example, allows an engineer working on safety requirements to check if the knowledge in Example 1 is consistent, or whether the system is safe on some particular, say 27th, day.

What if the engineer’s task is more general, as they need to compute the frequency with which the system from Example 1 is guaranteed to be safe and to determine on which days the system may be vulnerable? In this case the engineer needs a more in-depth understanding of the logical consequences of the given temporal knowledge, which can be provided by presenting all of the (infinitely many) temporal facts entailed by Formulas (1)–(4); note that this cannot be obtained using standard services like consistency checking or query answering. Computing all of the entailed temporal facts is both conceptually and computationally challenging. Hence, instead of asking the engineer to come up with a case-based solution, we would prefer to provide them

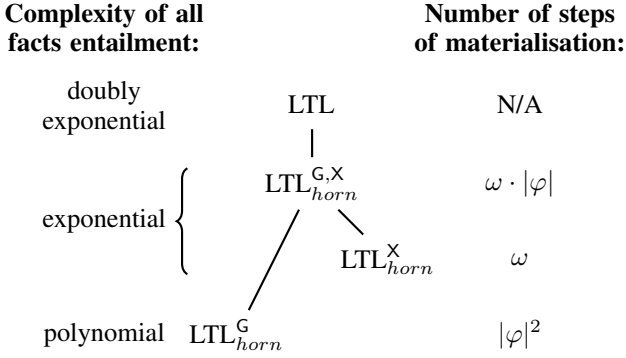


Figure 1: Cumulative results from this paper

with a general system performing reasoning automatically.

We address this task by developing optimal algorithms for computing all facts entailed by an LTL-formula. With the reasoning procedures developed in this paper, an engineer can effortlessly determine that the system from Example 1 is only guaranteed to be safe on the days of the form $2520 \cdot n$, for all $n \in \mathbb{N}$. Hence they may want to consider taking additional measures to increase the safety of the system.

Our contributions in this paper are as follows:

- After presenting preliminaries in Section 2, in Section 3 we define our problem formally as computing the intersection of all traces satisfying a given LTL-formula.
- In Section 4 we show that the problem can be solved in doubly exponential time which is worst-case unavoidable. We present an optimal algorithm for solving the problem which is based on manipulating a non-deterministic Büchi automaton corresponding to the input formula (which can be constructed using off-the-shelf systems).
- The negative complexity result from Section 4 motivates us to search for fragments of LTL in which our computational problem is easier. To this end, we study Horn fragments of LTL that were exploited in the setting of temporal ontology-based data access (Artale et al. 2013). In particular, we focus on Horn fragments $\text{LTL}_{\text{horn}}^{\text{G,X}}$, $\text{LTL}_{\text{horn}}^{\text{X}}$, and $\text{LTL}_{\text{horn}}^{\text{G}}$ which restrict temporal operators allowed in rules to the ones listed in the upper indices; the Hasse diagram of these fragments is depicted in Figure 1.
- In contrast to full LTL, each (satisfiable) formula φ in the fragments we consider has a unique least model, which can be computed via *materialisation*, that is, successive applications of Horn clauses seen as ‘if-then’ rules. As we show in Section 5, materialisation can take up to $\omega \cdot |\varphi|$ steps in $\text{LTL}_{\text{horn}}^{\text{G,X}}$, up to ω steps in $\text{LTL}_{\text{horn}}^{\text{X}}$, but only up to $|\varphi|^2$ steps in $\text{LTL}_{\text{horn}}^{\text{G}}$, as shown in Figure 1.
- In Section 6 we use the results on materialisation in Horn fragments to show that computing all facts entailed by $\text{LTL}_{\text{horn}}^{\text{G,X}}$ - and $\text{LTL}_{\text{horn}}^{\text{X}}$ -formulas requires exponential time. In the case of $\text{LTL}_{\text{horn}}^{\text{G}}$ -formulas, however, the problem is tractable, in particular, it becomes solvable in time $O(|\varphi|^4)$ —see Figure 1. Our algorithm for these frag-

ments is based on performing materialisation steps, where termination is ensured by applying specific stopping conditions and completeness of the construction is obtained as a result of *unfolding* partial materialisations.

- In Section 7 we present a prototype implementation of the algorithms from Sections 4 and 6. We perform experiments for full LTL and Horn fragments on randomly generated formulas. Materialisation-based approach vastly outperforms the automata-based approach, suggesting practical usefulness of the former.

2 Preliminaries

Linear Temporal Logic. Formulas of LTL (Pnueli 1977) are generated by the following grammar, where p is a proposition from a countable set PROP:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid X\varphi \mid \varphi U \varphi.$$

We use standard abbreviations such as $\top := p \vee \neg p$, for an arbitrary $p \in \text{PROP}$, $\perp := \neg\top$, $\varphi \wedge \psi := \neg(\neg\varphi \vee \neg\psi)$, $\varphi \rightarrow \psi := \neg\varphi \vee \psi$, $\varphi \leftrightarrow \psi := (\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$, $F\varphi := \top U \varphi$, $G\varphi := \neg F\neg\varphi$, and $X^k\varphi$ for φ preceded with $k \in \mathbb{N}$ operators X , where φ and ψ are any formulas. The *signature*, $\text{PROP}(\varphi)$, of a formula φ is the set of all propositions in φ . The *size* of φ , written $|\varphi|$, is the number of symbols in φ .

Formulas of LTL are interpreted over (infinite) *traces* $\sigma \in (2^{\text{PROP}})^\omega$; we also exploit *finite traces* $\sigma \in (2^{\text{PROP}})^*$ (we will always explicitly mention their finiteness) and denote their lengths by $|\sigma|$. *Satisfaction* of a formula φ at a time point $i \in \mathbb{N}$ of an infinite trace σ , denoted by $\sigma, i \models \varphi$, is defined below, where $\sigma(i)$ is the i th element of σ :

$$\begin{aligned} \sigma, i \models p & \quad \text{iff} \quad p \in \sigma(i) \\ \sigma, i \models \neg\varphi & \quad \text{iff} \quad \sigma, i \not\models \varphi \\ \sigma, i \models \varphi_1 \vee \varphi_2 & \quad \text{iff} \quad \sigma, i \models \varphi_1 \text{ or } \sigma, i \models \varphi_2 \\ \sigma, i \models X\varphi & \quad \text{iff} \quad \sigma, i+1 \models \varphi \\ \sigma, i \models \varphi_1 U \varphi_2 & \quad \text{iff} \quad \sigma, j \models \varphi_2 \text{ for some } j \geq i \text{ and} \\ & \quad \sigma, k \models \varphi_1 \text{ for all } k \text{ with } i \leq k < j \end{aligned}$$

A formula φ is *satisfied* by a trace σ if $\sigma, 0 \models \varphi$ and φ is *satisfiable* if it is satisfied by some trace. A formula φ *entails* a formula ψ if $\sigma, 0 \models \varphi$ implies $\sigma, 0 \models \psi$, for each trace σ . For $\sigma = a_0 a_1 \dots$ and any $i, j \in \mathbb{N}$, we let $\sigma(i) = a_i$ and $\sigma[i, j] := a_i a_{i+1} \dots a_j$. A trace σ is *contained* in a trace σ' , written as $\sigma \subseteq \sigma'$, if $\sigma(i) \subseteq \sigma'(i)$ for each $i \in \mathbb{N}$. The *intersection* $\bigcap \text{Tr}$ and the *union* $\bigcup \text{Tr}$ of traces from a set Tr are traces such that $\bigcap \text{Tr}(i) = \bigcap_{\sigma \in \text{Tr}} \sigma(i)$ and $\bigcup \text{Tr}(i) = \bigcup_{\sigma \in \text{Tr}} \sigma(i)$ for each $i \in \mathbb{N}$.

Büchi Automata. A (non-deterministic) *Büchi automaton* (NBA) is a tuple $\mathcal{A} = (\mathcal{Q}, \Sigma, \delta, q_0, \mathcal{F})$, where \mathcal{Q} is a finite set of *states*, Σ is a finite *alphabet*, $\delta : \mathcal{Q} \times \Sigma \rightarrow 2^{\mathcal{Q}}$ is a *transition function*, $q_0 \in \mathcal{Q}$ is an *initial state*, and $\mathcal{F} \subseteq \mathcal{Q}$ is a set of *final states*. An NBA \mathcal{A} accepts an ω -word $a_0 a_1 \dots \in \Sigma^\omega$ if there is a sequence $q_0, q_1, \dots \in \mathcal{Q}^\omega$ of states, called an *accepting run*, such that $q_{i+1} \in \delta(q_i, a_i)$ for each $i \in \mathbb{N}$ and some state in \mathcal{F} occurs infinitely often in q_0, q_1, \dots . The language of \mathcal{A} is the set $\mathcal{L}(\mathcal{A})$ of all ω -words accepted by

\mathcal{A} ; checking emptiness of $\mathcal{L}(\mathcal{A})$ is NL-complete (Vardi and Wolper 1994). For each LTL-formula φ there is an NBA \mathcal{A}_φ accepting exactly those infinite traces which satisfy φ ; the size of such \mathcal{A}_φ (i.e., the number of its states) is exponentially large in the size of φ (Vardi and Wolper 1994).

Horn Fragment of LTL. Following the approach of Artale et al. (2013)¹ we define Horn fragments of LTL by exploiting the clausal form of LTL-formulas, which was inspired by the *separated normal form* by Fisher (1991). To this end, we let a *temporal atom* be given by the following grammar, where $p \in \text{PROP}$:

$$\lambda ::= \perp \mid p \mid X\lambda \mid G\lambda \mid \lambda U \lambda.$$

An LTL-formula in *clausal form* is given by the grammar:

$$\varphi ::= X^k p \mid G(\neg\lambda_1 \vee \dots \vee \neg\lambda_n \vee \lambda_{n+1} \vee \dots \vee \lambda_{n+m}) \mid \varphi \wedge \varphi,$$

where n, m are positive integers, $k \in \mathbb{N}$, whereas λ and all λ_i are temporal atoms. Similar normal forms have been shown to preserve satisfiability of formulas (Artale et al. 2013; Fisher 1991); it is however not hard to see that they also preserve entailment, as we show below.

Proposition 2. *Each LTL-formula φ can be transformed in logarithmic space into φ' in clausal form, such that φ and φ' entail the same LTL-formulas over the signature of φ .*

Proof sketch. The transformation ‘renames’ subformulas ψ of φ to fresh propositions of the form a_ψ and introduces clauses which guarantee that each ψ and the corresponding a_ψ hold at the same time points. For example $\varphi = p \vee Xq$ is transformed into $\varphi' = a_\varphi \wedge G(a_\varphi \leftrightarrow a_p \vee a_{Xq}) \wedge G(a_{Xq} \leftrightarrow Xa_q) \wedge G(a_p \leftrightarrow p) \wedge G(a_q \leftrightarrow q)$, where an equivalence such as in $G(a_\varphi \leftrightarrow a_p \vee a_{Xq})$ is an abbreviation of $G(\neg a_\varphi \vee a_p \vee a_{Xq}) \wedge G(\neg a_p \vee a_\varphi) \wedge G(\neg a_{Xq} \vee a_\varphi)$. \square

Observe that a translation in the opposite direction to the one from Proposition 2 can be obtained by simply applying abbreviations from the beginning of this section, so that \perp , \wedge , and G are eliminated from a formula in clausal form.

A formula is *Horn* when it is in clausal form with $m = 1$. Such a formula can be seen as:

- a set of *facts* of the form $X^k p$, called a *dataset*, \mathcal{D} , and
- a set of *rules* of the form $G(\neg\lambda_1 \vee \dots \vee \neg\lambda_n \vee \lambda_{n+1})$, called a *program*, Π .

We let $|\mathcal{D}|$ and $|\Pi|$ be the number of symbols in \mathcal{D} and Π , respectively. Moreover, following the convention of logic programming, we will write a rule, r , as

$$\lambda_{n+1} \leftarrow \lambda_1 \wedge \dots \wedge \lambda_n,$$

where $\lambda_1 \wedge \dots \wedge \lambda_n$ is the *body* of r , in symbols $\text{body}(r)$, with each λ_i , for $i \leq n$, being its *body atom*, and λ_{n+1} is its *head*. If the head of r mentions \perp , we call r a \perp -*rule*. Similarly to Artale et al. (2013), we will pay special attention to Horn

formulas which allow only for G and X among temporal operators in rules, or only for X , or only for G . The corresponding classes: $\text{LTL}_{\text{horn}}^{G,X}$, $\text{LTL}_{\text{horn}}^X$, and $\text{LTL}_{\text{horn}}^G$ form a Hasse diagram depicted in Figure 1; note that these restrictions influence programs only, so the definition of datasets remains the same. It is also worth noting that there exists a number of other, yet similar, definitions of Horn temporal formulas in the literature (Chen and Lin 1993; Orgun and Ma 2005; Gabbay 1987; Abadi and Manna 1989).

Materialisation. *Materialisation* is a procedure that can be used for reasoning in $\text{LTL}_{\text{horn}}^{G,X}$, $\text{LTL}_{\text{horn}}^X$, and $\text{LTL}_{\text{horn}}^G$ (where rules do not mention disjunctions or U in their heads, and so, are ‘deterministic’). Given such a formula composed of a dataset \mathcal{D} and a program Π , *materialisation* consists in successively applying the *immediate consequence operator* T_Π to the least (with respect to containment) trace $\sigma_{\mathcal{D}}$ satisfying \mathcal{D} . The operator T_Π maps a trace σ to the least trace σ' containing σ and satisfying the following property for each $r \in \Pi$ which is not a \perp -rule: whenever σ satisfies each body atom of r at a time point i , then σ' satisfies the head of r at i ; by the form of $\text{LTL}_{\text{horn}}^{G,X}$, $\text{LTL}_{\text{horn}}^X$, and $\text{LTL}_{\text{horn}}^G$ -formulas, there exists exactly one such σ' , and so, $T_\Pi(\sigma)$ is well defined. Successive applications of T_Π to $\sigma_{\mathcal{D}}$ constitute a transfinite sequence of traces $T_\Pi^\alpha(\sigma_{\mathcal{D}})$ for ordinals α :

$$\begin{aligned} T_\Pi^0(\sigma_{\mathcal{D}}) &= \sigma_{\mathcal{D}}, \\ T_\Pi^{\alpha+1}(\sigma_{\mathcal{D}}) &= T_\Pi(T_\Pi^\alpha(\sigma_{\mathcal{D}})), \quad \text{for } \alpha \text{ an ordinal,} \\ T_\Pi^\alpha(\sigma_{\mathcal{D}}) &= \bigcup_{\beta < \alpha} T_\Pi^\beta(\sigma_{\mathcal{D}}), \quad \text{for } \alpha \text{ a limit ordinal.} \end{aligned}$$

We let $\mathfrak{C}_{\Pi, \mathcal{D}} = T_\Pi^{\omega_1}(\sigma_{\mathcal{D}})$, where ω_1 is the first uncountable ordinal; if $\mathfrak{C}_{\Pi, \mathcal{D}}$ satisfies all \perp -rules in Π , then $\mathfrak{C}_{\Pi, \mathcal{D}}$ is the least trace satisfying both Π and \mathcal{D} (Brandt et al. 2017). We say that materialisation of Π and \mathcal{D} takes α steps if α is the least ordinal such that $T_\Pi^\alpha(\sigma_{\mathcal{D}}) = \mathfrak{C}_{\Pi, \mathcal{D}}$.

Complexity of Reasoning. Satisfiability of LTL-formulas over infinite traces (Sistla and Clarke 1985) as well as over finite traces (De Giacomo and Vardi 2013) is PSpace-complete. Satisfiability of Horn formulas over \mathbb{Z} in fragments similar to our $\text{LTL}_{\text{horn}}^{G,X}$ and $\text{LTL}_{\text{horn}}^G$ (but with past temporal operators and over the irreflexive semantics) is PSpace- and P-complete, respectively (Artale et al. 2013). More recently, a number of problems in Horn temporal logics have been studied for their complexity, including satisfiability in interval logics (Bresolin et al. 2017; Wałęga 2023), rewritability of ontology-mediated queries with LTL and metric temporal logic (MTL) operators (Artale et al. 2021; Ryzhikov, Wałęga, and Zakharyashev 2020), as well as fact entailment in extensions of Datalog with MTL-operators (Wałęga et al. 2019; Wałęga et al. 2020; ?).

3 Intersection of All Traces

The main problem we consider in this paper is to efficiently compute all the atomic facts entailed by a given LTL-formula φ . One way to address this problem is to consider

¹In contrast to Artale et al. (2013), however, we use the standard syntax of LTL with no past operators, we adapt the standard semantics of LTL which is over \mathbb{N} and not over \mathbb{Z} , and we exploit the ‘reflexive’ rather than ‘irreflexive’ semantics of U and G .

all the traces satisfying φ and to check, for every position i , which atoms p occur in all these traces at this position. This would allow us to detect, for each i , all the atoms entailed by φ at i . In other words, this approach consists in computing the intersection of all traces satisfying φ . This intersection is a trace itself, denoted by σ_φ , and defined formally below.

Definition 3. Let φ be an LTL-formula and let Tr_φ be the set of all traces satisfying φ at the time point 0. We define the intersection of all traces of φ as the trace $\sigma_\varphi = \bigcap \text{Tr}_\varphi$.

In the definition above we adapt the convention that if $\text{Tr}_\varphi = \emptyset$, then $\bigcap \text{Tr}_\varphi$ is the ‘full trace’ over the signature of φ , $(\text{PROP}(\varphi))^\omega$. Thus, σ_φ is well-defined even if φ is unsatisfiable. At the same time it is worth noting that it is not guaranteed that σ_φ satisfies φ . For example if $\varphi = \text{F}p$, then σ_φ is an empty trace \emptyset^ω , and so, $\sigma_\varphi, 0 \not\models \varphi$. On the other hand, if φ is an $\text{LTL}_{\text{horn}}^{\text{G,X}}$ -formula (which clearly also applies to more restricted $\text{LTL}_{\text{horn}}^{\text{X}}$ - and $\text{LTL}_{\text{horn}}^{\text{G}}$ -formulas) then, as we described in Section 2, if φ is satisfiable, it has the unique least trace. This least trace is the intersection of all traces of φ , and so, it coincides with σ_φ . Therefore, in contrast to the general case, we obtain that $\sigma_\varphi, 0 \models \varphi$.

The problem we address in this paper is to efficiently compute σ_φ . Since σ_φ is an infinite trace, we start by providing a way to represent it finitely. To this end, we observe, as stated below, that σ_φ needs to be an *ultimately-periodic* trace.

Proposition 4. For any LTL-formula φ we have $\sigma_\varphi = uw^\omega$, where u and w are some finite traces.

Note that it is well known that each satisfiable LTL-formula has an ultimately-periodic trace (this fact is heavily exploited in LTL reasoners), but to show that σ_φ is also ultimately periodic, additional argumentation is in place. For that purpose, we can construct a formula of S1S (monadic second-order logic of one successor) containing

$$\forall x(p_i(x) \leftrightarrow \forall X_1, \dots, X_n(\text{tr}_\varphi(X_1, \dots, X_n) \rightarrow x \in X_i)),$$

for each of n propositions p_i in the signature of φ . It states that for each position x , the proposition p_i holds at x if and only if p_i holds at x in all traces of φ . For this, we construct a formula $\text{tr}_\varphi(X_1, \dots, X_n)$ with set variables X_1, \dots, X_n , which expresses that φ holds at 0 when, for each proposition p_i in φ , p_i holds at every position from X_i and only there. Any S1S formula can be transformed into a Büchi automaton and each Büchi automaton accepting some word accepts an ultimately periodic word. In our case, the automaton accepts exactly one word σ_φ , so that σ_φ is ultimately periodic.

Next, we observe that since σ_φ is ultimately periodic, there need to exist unique u and w with the least sum of lengths which represent σ_φ . We will call such a pair (u, w) the shortest representation of σ_φ as stated formally below. It is also worth observing that the shortest representation of a periodic word can be computed from a non-shortest one in linear time (Crochemore and Rytter 1994).

Definition 5. For an LTL-formula φ we let the shortest representation of σ_φ be the pair (u, w) of finite traces with minimal sum of lengths $|u| + |w|$ such that $uw^\omega = \sigma_\varphi$.

Now, based on the above discussion on the form of σ_φ and its representation, we are ready to formulate the main com-

putational problem considered in this paper—computing the intersection of all traces satisfying a given LTL-formula.

Definition 6. We let the intersection of traces be the following computational problem:

Input: an LTL-formula φ ,

Output: the shortest representation (u, w) of σ_φ .

By the definition, the intersection of traces provides us a complete information about entailment of all propositions by a formula φ . Furthermore, we can use intersection of traces to check entailment of complex formulas by simulating them with fresh propositions; for example, to determine all the positions in which a complex formula ψ is entailed by φ , it suffices to compute the intersection of traces for the formula $\varphi' = \varphi \wedge \text{G}(p_\psi \leftrightarrow \psi)$, which simulates ψ with a fresh proposition p_ψ . Then, all the positions in which p_ψ holds in σ_φ coincide with the positions in which ψ is entailed by φ .

Finally, we note that the argument (via constructing an S1S formula) for Proposition 4 implies that the intersection of traces is computable. The involved construction of a Büchi automaton from an S1S-formula is, however, non-elementary. In the next sections we will show how to compute the intersection of traces efficiently.

4 An Optimal Algorithm for LTL

In this section, we present an optimal algorithm for computing the intersection of traces of an arbitrary LTL-formula φ . We will prove, however, that for full LTL the worst-case complexity of the problem is doubly exponential.

Algorithm 1: Intersection of traces for LTL

Input: an LTL-formula φ
Output: the shortest representation (u, w) of σ_φ if φ is satisfiable, or **unsat** otherwise

- 1 Construct an NBA $\mathcal{A}_\varphi = (\mathcal{Q}, \Sigma, \delta, q_0, F)$ for φ ;
- 2 $j := 0$; $\mathcal{Q}_0 := \{q_0\}$;
- 3 **loop**
- 4 $\mathcal{Q}_{j+1} := \{q' \in \mathcal{Q} \mid q' \in \delta(q, a) \text{ for some } q \in \mathcal{Q}_j, a \in \Sigma, \mathcal{A}_\varphi \text{ has an accepting run from } q'\}$;
- 5 **if** $\mathcal{Q}_{j+1} = \emptyset$ **then return unsat**;
- 6 $a_j := \bigcap \{a \in \Sigma \mid \delta(q, a) \in \mathcal{Q}_{j+1}, q \in \mathcal{Q}_j\}$;
- 7 **if** $\mathcal{Q}_i = \mathcal{Q}_j$ for some $i < j$ **then**
- 8 $(u, w) :=$ the shortest representation of $a_0 \cdots a_{i-1}(a_i \cdots a_{j-1})^\omega$;
- 9 **return** (u, w) ;
- 10 $j := j + 1$;

Algorithm 1 presents our approach. In Line 1 it constructs an NBA \mathcal{A}_φ accepting the infinite traces which satisfy φ (Vardi and Wolper 1994), where each symbol a in the alphabet $\Sigma = 2^{\text{PROP}(\varphi)}$ of \mathcal{A}_φ is a subset of propositions occurring in φ . In Line 2, it initialises counter j to 0 and \mathcal{Q}_0 to the singleton containing only the initial state q_0 of \mathcal{A}_φ . Then, each iteration of the loop in Lines 3–9: computes the set \mathcal{Q}_{j+1} of states which can be accessed from states in \mathcal{Q}_j and from which \mathcal{A}_φ has an accepting run (Line 4), computes $a_j = \sigma_\varphi(j)$ (Line 5), and increments j by 1 (Line 9).

If some \mathcal{Q}_j is empty, then φ is unsatisfiable (Line ??) and if some \mathcal{Q}_i and \mathcal{Q}_j have the same elements (Line 6), the trace $a_0 \cdots a_{i-1} (a_i \cdots a_{j-1})^\omega$ coincides with σ_φ . Finally, its shortest representation (u, w) is computed in Line 7 in linear time (Crochemore and Rytter 1994).

Theorem 7. *Algorithm 1 returns the shortest representation of σ_φ , or *unsat* if φ is unsatisfiable; the algorithm terminates in time $2^{2^{O(|\varphi|)}}$.*

Proof. Note that φ is unsatisfiable if and only if \mathcal{Q}_1 is empty, so we can focus on satisfiable formulas.

Let (u, w) be the output of the algorithm, so it is computed in Line 7 as the shortest representation of $a_0 \cdots a_{i-1} (a_i \cdots a_{j-1})^\omega$. Hence, it suffices to show that the latter coincides with σ_φ . First, we show by induction on $k \in \mathbb{N}$ that if q_0, q_1, \dots is an accepting run of \mathcal{A}_φ , then $q_k \in Q_\ell$, where $\ell = k$ if $k < j$, and $\ell = i + ((k - i) \bmod (j - i))$ if $k \geq j$. The base case (for $k = 0$) holds trivially. For the inductive step, if $k < j$ then $q_{k-1} \in Q_{k-1}$ so, by Line 4, $q_k \in Q_k$. If $k \geq j$ then $k = i + ((k - i) \bmod (j - i)) + n \cdot (j - i)$, for some $n \in \mathbb{N}$, and $q_{k-1} \in Q_\ell$ for $\ell = i + ((k - 1 - i) \bmod (j - i))$. If $(k - i) \bmod (j - i) \neq 0$, then $q_k \in Q_{\ell+1}$. Otherwise $q_k \in Q_i$, which completes the inductive step.

By the property proved above, for any $\sigma \in L(\mathcal{A}_\varphi)$ and $k < j$, after reading $\sigma[0, k]$, the automaton \mathcal{A}_φ is in a state from the set Q_k . Since $\sigma_\varphi(k) = \bigcap_{\sigma \in L(\mathcal{A}_\varphi)} \sigma(k)$, Line 5 implies that $\sigma_\varphi(k) = a_k$ if $k < j$. For $k \geq j$ we need to show that $\sigma_\varphi(k) = a_\ell$, where $\ell = i + ((k - i) \bmod (j - i))$. This, however, follows from the fact that after reading $\sigma[0, k]$, the automaton \mathcal{A}_φ is in a state from the set Q_ℓ .

For the doubly exponential bound on the runtime, we observe that \mathcal{A}_φ is exponentially large (Vardi and Wolper 1994), namely $|\mathcal{Q}| = O(2^{|\varphi|})$, so it suffices to show that Algorithm 1 terminates in time $O(2^{|\mathcal{Q}|})$. This holds since $\mathcal{Q}_j \subseteq \mathcal{Q}$, for each \mathcal{Q}_j constructed by the algorithm (see Line 4), and so, the if-condition from Line 6 yields termination after at most $2^{|\mathcal{Q}|}$ iterations of the main loop. \square

We close this section by showing that Algorithm 1 is optimal, as calculating the intersection of traces for some LTL-formulas requires doubly exponential time.

Proposition 8. *There is a family of LTL-formulas φ_n , with $n \in \mathbb{N}$, such that each φ_n is of polynomial size in n and σ_{φ_n} has the shortest representation (u, w) with $|w| = 2^{2^{\Omega(n)}}$.*

Proof. We will use φ_n to encode a binary counter which is initialised with any value $m < 2^n$ and counts up to $2^n - 1$, then sets the counter to m and repeats such computations *ad infinitum*. Formula φ_n will use propositions $0, 1, \#, \triangleleft$ and each of its (minimal) satisfying traces will be of the form

$$(\text{bin}_n(m) \# \text{bin}_n(m+1) \# \cdots \# \text{bin}_n(2^n - 1) \triangleleft)^\omega,$$

where $\text{bin}_n(m)$ is the encoding of m using n bits $b_0 b_1 \cdots b_{n-1} \in \{0, 1\}^n$ with the least-significant bit first.

The intersection of the traces above, for all $0 \leq m \leq 2^n - 1$, is the trace $(\emptyset^{(n+1) \cdot \ell - 1} \triangleleft)^\omega$, with

$\ell = \text{lcm}\{1, \dots, 2^n\}$. As $\text{lcm}\{1, \dots, 2^n\} \geq 2^{2^n}$ (Nair 1982), we get the required doubly exponential bound.

It remains to show how to encode our specific counter with φ_n . To this end we let φ_n be the conjunction

$\text{Segment}_n \wedge \text{Invariant}_n \wedge \text{End}_n \wedge \text{Propagates}_n \wedge \text{Counter}_n$.

We start by defining Segment_n which will hold in the time points $0, n+1, 2(n+1), \dots$, corresponding to the positions of the least significant bits b_0 :

$$\text{Segment}_n = \bigwedge_{0 \leq i \leq n-1} X^i(0 \vee 1) \wedge X^n(\triangleleft \vee \#).$$

Formula Invariant_n ensures that each segment is immediately (i.e., after $n+1$ time points) followed by another one:

$$\text{Invariant}_n = G(\text{Segment}_n \rightarrow X^{n+1} \text{Segment}_n).$$

We use End_n to state that \triangleleft ends segments with only 1s:

$$\text{End}_n = G(X^n \triangleleft \leftrightarrow \bigwedge_{0 \leq i \leq n-1} X^i 1).$$

Formula Propagate_n ensures that each segment ended by \triangleleft is followed by a segment encoding the initial value m :

$$\text{Propagate}_n = \bigwedge_{0 \leq i \leq n-1} \bigwedge_{j \in \{0,1\}} (X^i j \rightarrow G(\triangleleft \rightarrow X^{i+1} j)).$$

Finally, we use Counter_n to increment the value of the binary counter in the next segment. To this end, we introduce additional formulas $\text{Equal}_{n,i}$ and $\text{Inc}_{n,i}$ with $0 \leq i \leq n-1$ (for definiteness we additionally let $\text{Inc}_{n,n} = \top$):

$$\begin{aligned} \text{Equal}_{n,i} &= \bigwedge_{i \leq j \leq n-1} X^i 1 \leftrightarrow X^{i+n+1} 1, \\ \text{Inc}_{n,i} &= (X^i 0 \rightarrow (X^{i+n+1} 1 \wedge \text{Equal}_{n,i+1})) \wedge \\ &\quad (X^i 1 \rightarrow (X^{i+n+1} 0 \wedge \text{Inc}_{n,i+1})), \\ \text{Counter}_n &= G(X^n \# \rightarrow \text{Inc}_{n,0}). \end{aligned}$$

The size of φ_n is polynomial in n , as required. \square

In the next sections we address the negative result from Proposition 8 by showing Horn fragments of LTL in which computing the intersection of traces is significantly easier.

5 Materialisation in Horn Fragments

In this section, we study Horn fragments of LTL, where materialisation can be used to check satisfiability of formulas and, as we will show in Section 6, to efficiently compute the intersection of traces. In the first step of our study on Horn fragments, we determine the number of steps materialisation takes before reaching a fixpoint, which is crucial for both of the above-mentioned reasoning tasks. In particular, we will show that materialisation in $\text{LTL}_{\text{horn}}^G$, $\text{LTL}_{\text{horn}}^X$, and $\text{LTL}_{\text{horn}}^{X,G}$, takes up to $|\varphi|^2$, ω , and $\omega \cdot |\varphi|$ steps, respectively.

First we show the quadratic bound for $\text{LTL}_{\text{horn}}^G$ -formulas; we recall that $\sigma_{\mathcal{D}}$ is the least trace satisfying \mathcal{D} .

Theorem 9. *For every $\text{LTL}_{\text{horn}}^G$ -formula φ , composed of a program Π and dataset \mathcal{D} , we have that $\mathfrak{C}_{\Pi, \mathcal{D}} = T_{\Pi}^{|\varphi|^2}(\sigma_{\mathcal{D}})$.*

Proof sketch. We start by dividing the time line into segments $[0, k_1), [k_1, k_1], (k_1, k_2), \dots, [k_\ell, k_\ell], (k_\ell, \infty)$, where $k_1 < \dots < k_\ell$ are all the time points in \mathcal{D} (i.e., the numbers k from expressions $X^k p$ in \mathcal{D}). Since Π mentions only G among temporal operators, we can show by transfinite induction on ordinals α that $T_\Pi^\alpha(\sigma_{\mathcal{D}})$ satisfies the same propositions in all time points belonging to the same segment and that $T_\Pi^\alpha(\sigma_{\mathcal{D}})$ does not satisfy any proposition in the segment $[0, k_1)$. Thus each application of T_Π (before reaching $\mathfrak{C}_{\Pi, \mathcal{D}}$) derives at least one proposition in one of the $\leq 2 \cdot |\mathcal{D}|$ segments. There are at most $\frac{|\Pi|}{3}$ propositions that can be derived (each proposition is derived by some rule and each rule has at least 3 symbols), so $\mathfrak{C}_{\Pi, \mathcal{D}} = T_\Pi^\alpha(\sigma_{\mathcal{D}})$, for $\alpha = 2 \cdot |\mathcal{D}| \cdot \frac{|\Pi|}{3}$, and therefore, $\alpha \leq |\varphi|^2$. \square

Moreover, we can show that the bound from Theorem 9 is optimal by constructing a formula for which materialisation takes a quadratic number of steps. Indeed, consider a dataset \mathcal{D} consisting of $X^1 r_1$ and $X^i e_i$ for each $i \in \{1, \dots, k\}$, as well as a program Π with rules:

$$\begin{aligned} p_1 &\leftarrow e_i \wedge r_i, & \text{for all } i \in \{1, \dots, k\}, \\ p_{i+1} &\leftarrow p_i, & \text{for all } i \in \{1, \dots, \ell - 1\}, \\ G r_{i+1} &\leftarrow p_\ell \wedge e_i, & \text{for all } i \in \{1, \dots, k - 1\}. \end{aligned}$$

Materialisation of Π and \mathcal{D} takes $k \cdot (\ell + 1)$ steps. The first $\ell + 1$ steps derive facts at 1; the first step derives p_1 at 1 (using the first rule of Π), the next $\ell - 1$ steps derive p_2, \dots, p_ℓ at 1 (the second rule), and then $G r_2$ is derived at 1 (the third rule). The next $\ell + 1$ steps derive propositions at 2, and this process repeats k times, giving a quadratic number of steps.

In the remaining part of this section we show bounds on the number of materialisation steps for LTL_{horn}^X and $LTL_{horn}^{X,G}$. For this, the following notion will be useful.

Definition 10. For an $LTL_{horn}^{G,X}$ program Π , dataset \mathcal{D} , and an ordinal α , we let $\text{rules}_{\Pi, \mathcal{D}}^\alpha$ be the set of all rules $r \in \Pi$ for which there is a time point i such that $T_\Pi^\alpha(\sigma_{\mathcal{D}}), i \models \text{body}(r)$, but $T_\Pi^\beta(\sigma_{\mathcal{D}}), i \not\models \text{body}(r)$ for all $\beta < \alpha$.

Intuitively $\text{rules}_{\Pi, \mathcal{D}}^\alpha$ contains all rules in Π which ‘fire’ in step α of materialisation at some time point i , and which did not ‘fire’ at this i in any previous step. We show next that if α is a limit ordinal, all such rules mention G in the body.

Lemma 11. If α is a limit ordinal, then each rule in $\text{rules}_{\Pi, \mathcal{D}}^\alpha$ mentions G in the body, where Π is any $LTL_{horn}^{G,X}$ program and \mathcal{D} is any dataset.

Proof. Suppose towards a contradiction that there exists $r \in \text{rules}_{\Pi, \mathcal{D}}^\alpha$ which does not mention G in the body. Since $r \in \text{rules}_{\Pi, \mathcal{D}}^\alpha$, by Definition 10 there exists i such that $T_\Pi^\alpha(\sigma_{\mathcal{D}}), i \models \text{body}(r)$, but $T_\Pi^\beta(\sigma_{\mathcal{D}}), i \not\models \text{body}(r)$, for all ordinals $\beta < \alpha$. As r does not mention G in the body, there is a unique minimal (finite) dataset \mathcal{D}_r such that $\sigma_{\mathcal{D}_r} \subseteq T_\Pi^\alpha(\sigma_{\mathcal{D}})$ and $\sigma_{\mathcal{D}_r}, i \models \text{body}(r)$. Since α is a limit ordinal, by the definition (see Section 2), $T_\Pi^\alpha(\sigma_{\mathcal{D}}) = \bigcup_{\beta < \alpha} T_\Pi^\beta(\sigma_{\mathcal{D}})$, and so, for each $X^k p \in \mathcal{D}_r$ there is an ordinal $\beta < \alpha$ such that $T_\Pi^\beta(\sigma_{\mathcal{D}}) \models X^k p$. Let γ be the largest among these

β , so $\gamma < \alpha$. Moreover, as for each ordinal β we have $T_\Pi^\beta(\sigma_{\mathcal{D}}) \subseteq T_\Pi^{\beta+1}(\sigma_{\mathcal{D}})$, we obtain that $T_\Pi^\gamma(\sigma_{\mathcal{D}}) \models \mathcal{D}_r$. Hence, $T_\Pi^\gamma(\sigma_{\mathcal{D}}), i \models \text{body}(r)$, raising a contradiction. \square

Next, we use the above lemma to show a bound on the number of materialisation steps in $LTL_{horn}^{X,G}$ depending on the number k of rules which mention G in the body. Note that LTL_{horn}^X is a special case of $LTL_{horn}^{X,G}$ with $k = 0$, so the result below provides also a bound for LTL_{horn}^X .

Theorem 12. For every $LTL_{horn}^{G,X}$ program Π and every dataset \mathcal{D} , we have $\mathfrak{C}_{\Pi, \mathcal{D}} = T_\Pi^{\omega \cdot (k+1)}(\sigma_{\mathcal{D}})$, where k is the number of rules in Π which mention G in their bodies.

Proof sketch. To show that $\mathfrak{C}_{\Pi, \mathcal{D}} = T_\Pi^{\omega \cdot (k+1)}(\sigma_{\mathcal{D}})$, it suffices to prove that $\text{rules}_{\Pi, \mathcal{D}}^{\omega \cdot (k+1)} = \emptyset$. To this end, we will show that for any limit ordinal α and ordinal $\beta < \alpha$, we have $\text{rules}_{\Pi, \mathcal{D}}^\alpha \cap \text{rules}_{\Pi, \mathcal{D}}^\beta = \emptyset$. This will finish the proof as, by Lemma 11, for each limit ordinal α the set $\text{rules}_{\Pi, \mathcal{D}}^\alpha$ contains only rules with G in the body, and so, our result above implies that there are at most k limit ordinals α for which $\text{rules}_{\Pi, \mathcal{D}}^\alpha \neq \emptyset$. Thus, $\text{rules}_{\Pi, \mathcal{D}}^{\omega \cdot (k+1)} = \emptyset$.

To show the missing part of the proof, we suppose towards a contradiction that for some limit ordinal α and an ordinal $\beta < \alpha$, there is $r \in \text{rules}_{\Pi, \mathcal{D}}^\alpha \cap \text{rules}_{\Pi, \mathcal{D}}^\beta$. Since $r \in \text{rules}_{\Pi, \mathcal{D}}^\alpha$, there is a time point i such that $T_\Pi^\alpha(\sigma_{\mathcal{D}}), i \models \text{body}(r)$ but $T_\Pi^\beta(\sigma_{\mathcal{D}}), i \not\models \text{body}(r)$. As $r \in \text{rules}_{\Pi, \mathcal{D}}^\beta$, there is also a time point j such that $T_\Pi^\beta(\sigma_{\mathcal{D}}), j \models \text{body}(r)$, and clearly $i \neq j$. Now, by the fact that $T_\Pi^\beta(\sigma_{\mathcal{D}}) \subseteq T_\Pi^\alpha(\sigma_{\mathcal{D}})$, we can show that there exists a unique minimal (finite) dataset \mathcal{D}_r such that $\sigma_{\mathcal{D}_r} \subseteq T_\Pi^\alpha(\sigma_{\mathcal{D}})$ and $\sigma_{\mathcal{D}_r} \cup T_\Pi^\beta(\sigma_{\mathcal{D}}), i \models \text{body}(r)$. Next, by the same argument as in the proof of Lemma 11, we can use the fact that $T_\Pi^\alpha(\sigma_{\mathcal{D}}) = \bigcup_{\gamma < \alpha} T_\Pi^\gamma(\sigma_{\mathcal{D}})$ to show that there exists $\gamma < \alpha$ such that $T_\Pi^\gamma(\sigma_{\mathcal{D}}), i \models \text{body}(r)$. As $\gamma < \alpha$, it contradicts the assumption that $r \in \text{rules}_{\Pi, \mathcal{D}}^\alpha$. \square

Observe that we can replace $k + 1$ in Theorem 12 with $|\varphi|$. Indeed, $k + 1 \leq |\varphi|$ unless $|\varphi| = 0$, but in this border case case materialisation takes 0 steps. Note also that in LTL_{horn}^X we have $k = 0$, thus materialisation takes up to ω steps.

To show that these bounds are optimal, for any k we can construct Π and \mathcal{D} whose materialisation takes $\omega \cdot (k + 1)$ steps. Indeed, let $\mathcal{D} = \{X^0 p, X^0 q_1\}$ and let Π have rules:

$$\begin{aligned} X q_i &\leftarrow q_i, & \text{for all } i \in \{1, \dots, k + 1\}, \\ q_{i+1} &\leftarrow p \wedge G q_i, & \text{for all } i \in \{1, \dots, k\}. \end{aligned}$$

After ω applications of the first rule we get $G q_1$ at 0, allowing us to apply the second rule and to get q_2 at 0. This repeats $k + 1$ times giving rise to $\omega \cdot (k + 1)$ steps of materialisation.

We summarise the results from this sections below.

Corollary 13. Let φ be an $LTL_{horn}^{G,X}$ -formula consisting of a program Π and a dataset \mathcal{D} . Then, the following hold:

$$\begin{aligned} \mathfrak{C}_{\Pi, \mathcal{D}} &= T_\Pi^{\omega \cdot |\varphi|}(\sigma_{\mathcal{D}}), & \text{if } \varphi \text{ is an } LTL_{horn}^{G,X}\text{-formula,} \\ \mathfrak{C}_{\Pi, \mathcal{D}} &= T_\Pi^\omega(\sigma_{\mathcal{D}}), & \text{if } \varphi \text{ is an } LTL_{horn}^X\text{-formula,} \\ \mathfrak{C}_{\Pi, \mathcal{D}} &= T_\Pi^{|\varphi|^2}(\sigma_{\mathcal{D}}), & \text{if } \varphi \text{ is an } LTL_{horn}^G\text{-formula.} \end{aligned}$$

Our results improve the previously established bound ω_1 (Brandt et al. 2017) and are crucial for our algorithm in the next section, which optimally computes the intersection of traces for LTL_{horn}^G , LTL_{horn}^X , and $LTL_{horn}^{G,X}$ -formulas.

6 Algorithm for Horn Fragments

We will use results from Section 5 to compute the intersection of traces for formulas of $LTL_{horn}^{G,X}$ (as well as of LTL_{horn}^G and LTL_{horn}^X) or return `unsat` if the formula is unsatisfiable. Our algorithm takes as an input a formula in *normal form* in which the program has only rules of the forms:

$$Xq \leftarrow p, \quad Gq \leftarrow p, \quad q \leftarrow Xp, \quad q \leftarrow Gp, \quad r \leftarrow p \wedge q,$$

for $p, q \in \text{PROP}$ and $r \in \text{PROP} \cup \{\perp\}$. Each $LTL_{horn}^{G,X}$ -formula φ can be transformed in logarithmic space into a formula φ' in normal form such that $\sigma_{\varphi'}$ is a conservative extension of σ_{φ} , so the shortest representation (u, w) of σ_{φ} can be easily computed from the shortest representation (u', w') of $\sigma_{\varphi'}$. To construct φ' , we introduce rules with fresh propositions, for example, $XGq \leftarrow p \wedge q \wedge r$ is translated into $p_1 \leftarrow p \wedge q, p_2 \leftarrow p_1 \wedge r, Xp_3 \leftarrow p_2$, and $Gq \leftarrow p_3$; our construction guarantees also that φ' is in the same class ($LTL_{horn}^{G,X}$, LTL_{horn}^G , or LTL_{horn}^X) as φ .

Given a formula φ in normal form, which consists of a program Π and a dataset \mathcal{D} , Algorithm 2 shows how to efficiently compute the shortest representation of σ_{φ} . In each iteration of the main loop (Lines 1–9), the algorithm keeps extending \mathcal{D} . In particular, \mathcal{D} is extended in Line 4 using $\text{ApplyRules}(\Pi, \mathcal{D})$ which implements a one-step-application of T_{Π} to $\sigma_{\mathcal{D}}$. If some \perp -rules ‘fires’, the input formula is unsatisfiable, so the algorithm returns `unsat` in Line ???. This, however, does not prevent from applying rules infinitely many times. Hence we keep applying them only until \mathcal{D} enjoys a specific *saturation* property (see Definition 14), which allows us to compute the *period* $[i, j]$ of \mathcal{D} (Line 7). Next, in Line 8, the crucial part of the algorithm occurs, where \mathcal{D} and $[i, j]$ are used to compute an additional set \mathcal{D}_{∞} of formulas which follow from Π and \mathcal{D} , but which may not be derivable in any finite number of ApplyRules applications. After adding \mathcal{D}_{∞} to \mathcal{D} (Line 9) we may be able to derive further formulas, and so, we repeat the main loop until \mathcal{D}_{∞} is empty. Once the loop terminates, we can show that σ_{φ} coincides with the trace $\sigma_{\mathcal{D}}[0, i-1] (\sigma_{\mathcal{D}}[i, j])^{\omega}$, so it remains to compute its shortest representation (Line 11).

The crucial part of the algorithm is the definition of a (Π, k_{\max}) -saturated set and its period, which are used to compute \mathcal{D}_{∞} .

Definition 14. Let \mathcal{D} be a set of formulas of the forms $X^k p$ and $X^k Gp$, let Π be a program, and let k_{\max} be a time point. We say that \mathcal{D} is (Π, k_{\max}) -saturated if there is an interval $[i, j]$ with $k_{\max} < i$ such that

- $\sigma_{\mathcal{D}}, i \models \lambda$ if and only if $\sigma_{\mathcal{D}}, j+1 \models \lambda$, for each temporal atom λ occurring in Π , and
- $\sigma_{\mathcal{D}}, k \models \lambda_1 \wedge \dots \wedge \lambda_n \rightarrow \lambda_{n+1}$, for each $k \in [0, j]$ and each rule $\lambda_{n+1} \leftarrow \lambda_1 \wedge \dots \wedge \lambda_n$ in Π .

We let the period of \mathcal{D} be the interval with the least j among $[i, j]$ ’s satisfying the properties above.

Algorithm 2: Intersection of traces for $LTL_{horn}^{G,X}$

Input: an $LTL_{horn}^{G,X}$ -formula φ in normal form consisting of a program Π and a dataset \mathcal{D}
Output: the shortest representation (u, w) of σ_{φ} if φ is satisfiable, or `unsat` otherwise

```

1 repeat
2    $k_{\max} :=$  maximal  $k$  such that  $X^k$  occurs in  $\mathcal{D}$ ;
3   while  $\mathcal{D}$  is not  $(\Pi, k_{\max})$ -saturated do
4     if  $\sigma_{\mathcal{D}}, i \models \text{body}(r)$ , for some  $i \in \mathbb{N}$  and a
        $\perp$ -rule  $r \in \Pi$  then return unsat;
5     else  $\mathcal{D} := \text{ApplyRules}(\Pi, \mathcal{D})$ ;
6    $[i, j] :=$  period of  $\mathcal{D}$ ;
7    $\mathcal{D}_{\infty} := \{X^i Gp \mid Gp \text{ occurs in the body of } \Pi,$ 
        $\sigma_{\mathcal{D}}, k \models p \text{ for all } k \in [i, j], \text{ and } \sigma_{\mathcal{D}}, i \not\models Gp\}$ ;
8    $\mathcal{D} := \mathcal{D} \cup \mathcal{D}_{\infty}$ ;
9 until  $\mathcal{D}_{\infty} = \emptyset$ ;
10  $(u, w) :=$  the shortest representation of
        $\sigma_{\mathcal{D}}[0, i-1] (\sigma_{\mathcal{D}}[i, j])^{\omega}$ ;
11 return  $(u, w)$ ;
```

Note that the period in Definition 14 is uniquely defined; indeed, if there were two distinct periods $[i, j]$ and $[i', j]$ with $i < i'$, then, by the first item in the definition, the same temporal atoms λ from Π would be satisfied at i and i' . Hence, $[i, i' - 1]$, with $i' - 1 < j$, would satisfy the properties of a period, which contradicts the minimality of j .

In the next theorems we show that Algorithm 2 is correct and worst-case optimal. For any $s \in \{k_{\max}, i, j, \mathcal{D}_{\infty}\}$, and any iteration $\ell \in \mathbb{N}$ of the main loop, we will use s^{ℓ} for the value of s at the end of the ℓ th iteration. We also let \mathcal{D}_i^{ℓ} and \mathcal{D}_e^{ℓ} be the values of \mathcal{D} in the ℓ th iteration before and after executing the ‘while’ loop, respectively.

Theorem 15. Algorithm 2 terminates in time $O(2^{|\varphi|})$ in general and in time $O(|\varphi|^4)$ if φ is an LTL_{horn}^G -formula.

Proof. The most significant component in the computation is the number A of applications of ApplyRules (taking place in Line 4). Each such application adds to \mathcal{D} at least one new formula of the form $X^k p$ or $X^k Gp$, which did not follow from \mathcal{D} so far. Thus, $A \leq k_{\max}^{f+1} \cdot 2^{|\varphi|}$, where k_{\max}^f is the value of k_{\max} in the final iteration f of the main loop, whereas we let k_{\max}^{f+1} be the value as if it was computed after the final iteration. We will show that $k_{\max}^{f+1} = O(2^{|\varphi|})$. For this, we observe that $k_{\max}^{\ell+1} \leq k_{\max}^{\ell} + (k_{\max}^{\ell} + 2^{|\varphi|} + 1) \cdot 2^{|\varphi|}$ for each iteration ℓ , since $(k_{\max}^{\ell} + 2^{|\varphi|} + 1) \cdot 2^{|\varphi|}$ is the maximal number of applications of ApplyRules till saturation is reached, and each such application can increase the maximal number in \mathcal{D} by at most 1. We can find a constant $c \in \mathbb{N}$ such that, for any ℓ , we have $2^{|\varphi|} + 1 \leq c \cdot k_{\max}^{\ell}$, so $k_{\max}^{\ell+1} \leq k_{\max}^{\ell} \cdot (1 + 4c^2 k_{\max}^{\ell}) \leq 5(c^2 k_{\max}^{\ell})^2$. Since $k_{\max}^1 \leq |\varphi|$ we obtain that $k_{\max}^{f+1} = O(|\varphi|^{f+1})$. It remains to show that $f = O(|\varphi|)$. In particular, $f < k + 1$, where $k \leq |\varphi|$ is the number of body atoms Gp in Π which mention distinct propositions p . This is because there can be at most k distinct and non-empty sets \mathcal{D}_{∞} in the run; otherwise we

had $X^i Gp \in \mathcal{D}_\infty^\ell$ and $X^{i'} Gp \in \mathcal{D}_\infty^{\ell'}$ for some $\ell' > \ell$ and $i, i' \in \mathbb{N}$. However, $\ell' > \ell$ implies $i' \geq i$, and so $X^i Gp$ entails $X^{i'} Gp$ contradicting $X^{i'} Gp \in \mathcal{D}_\infty^{\ell'}$ (see Line 8).

We showed an exponential bound on k_{\max}^ℓ , so each \mathcal{D}_b^ℓ (and \mathcal{D}_e^ℓ) is also exponentially bounded. We observe that each application of `ApplyRules`(Π, \mathcal{D}) is polynomial in the size of Π and current \mathcal{D} , and so is checking saturation of \mathcal{D} (Line 3) and computing its period (Line 7). This allows us to conclude that the algorithm runs in $O(2^{|\varphi|})$ time.

If φ is an $\text{LTL}_{\text{horn}}^G$ -formula, there is only one iteration of the algorithm's main loop and, as we have shown in Corollary 13, the number of `ApplyRules` calls in this iteration is in $O(|\varphi|^2)$. Moreover, it follows from the proof of Theorem 9 that each $X^k p$ and $X^k Gp$ in \mathcal{D} has k that occurs in the initial \mathcal{D}_b^1 . Thus, there are linearly many such k 's and linearly many rules that can be applied, so each call of `ApplyRules` requires $O(|\varphi|^2)$ time. Each saturation check and computing the period (which are actually not needed in $\text{LTL}_{\text{horn}}^G$) is also in $O(|\varphi|^2)$. Thus, the overall runtime is $O(|\varphi|^4)$. \square

Theorem 16. *Algorithm 2 returns the shortest representation of σ_φ , or `unsat` if φ is unsatisfiable.*

Proof sketch. First, we assume that φ has no \perp -rules, so it is satisfiable and the algorithm returns some (u, w) . We will show that $\sigma_\varphi = uw^\omega$. For any iteration ℓ we let σ^ℓ be $\sigma_{\mathcal{D}_e^\ell}[0, i^\ell - 1](\sigma_{\mathcal{D}_e^\ell}[i^\ell, j^\ell])^\omega$, thus $uw^\omega = \sigma^f$, for f being the final iteration. Hence, we want to show that $\sigma_\varphi = \sigma^f$.

(\subseteq) As φ is a satisfiable $\text{LTL}_{\text{horn}}^{X,G}$ -formula, σ_φ is the least trace of φ . Hence, to prove that $\sigma_\varphi \subseteq \sigma^f$, it suffices to show that σ^f is a trace of φ . By the structure of the algorithm, no formula is ever removed from \mathcal{D} , so $\sigma^f, 0 \models \mathcal{D}$. Thus, it remains to prove that $\sigma^f, 0 \models \Pi$, that is, for each time point t and each rule $\lambda_3 \leftarrow \lambda_1 \wedge \lambda_2$ in Π , if $\sigma^f, t \models \lambda_1 \wedge \lambda_2$ then $\sigma^f, t \models \lambda_3$. We observe that $\sigma_{\mathcal{D}_e^f}$ satisfies this property for all $t \in [0, j^f]$, since it is (Π, k_{\max}^f) -saturated. Hence, it suffices to show that for each t and atom λ occurring in Π the following statements hold:

- (i) if $t < i^f$, then $\sigma^f, t \models \lambda$ iff $\sigma_{\mathcal{D}_e^f}, t \models \lambda$, and
- (ii) if $t \geq i^f$, then $\sigma^f, t \models \lambda$ iff $\sigma_{\mathcal{D}_e^f}, t' \models \lambda$, where t' is the unique point in $[i^f, j^f]$ such that $j^f + 1 - i^f$ divides $t - t'$.

As φ is in normal form, each λ is of the form p , Gp , or Xp , so we conduct a proof by cases. If $\lambda = p$, Statements (i) and (ii) hold since $\sigma^f = \sigma_{\mathcal{D}_e^f}[0, i^f - 1](\sigma_{\mathcal{D}_e^f}[i^f, j^f])^\omega$. If $\lambda = Xp$, Statement (i) holds for the same reason, and so does Statement (ii) for $t' \neq j^f$. For $t' = j^f$, the following are equivalent: (1) $\sigma^f, t \models Xp$, (2) $\sigma^f, t + 1 \models p$, (3) $\sigma_{\mathcal{D}_e^f}, i^f \models p$ (as $(t + 1)' = i^f$), (4) $\sigma_{\mathcal{D}_e^f}, j^f + 1 \models p$ (as $[i^f, j^f]$ is the period of \mathcal{D}_e^f), and (5) $\sigma_{\mathcal{D}_e^f}, j^f \models Xp$. If $\lambda = Gp$, the proof is similar, but uses the fact that $\mathcal{D}_\infty^f = \emptyset$.

(\supseteq) Next we prove that $\sigma^f \subseteq \sigma_\varphi$. We claim that it suffices to show that (*) for any $\ell \leq f$ if $\sigma_{\mathcal{D}_b^\ell} \subseteq \sigma_\varphi$, then $\sigma^\ell \subseteq \sigma_\varphi$. Statement (*) indeed implies that $\sigma^f \subseteq \sigma_\varphi$; since $\sigma_{\mathcal{D}_b^1} \subseteq \sigma_\varphi$ and, moreover, $\sigma^\ell \subseteq \sigma_\varphi$ implies that $\sigma_{\mathcal{D}_b^{\ell+1}} \subseteq \sigma_\varphi$ (by the

definition of \mathcal{D}_∞^ℓ , which is used to construct $\mathcal{D}_b^{\ell+1}$), by Statement (*) we obtain that $\sigma_{\mathcal{D}_b^f} \subseteq \sigma_\varphi$ and also $\sigma^f \subseteq \sigma_\varphi$.

To prove Statement (*), we fix $\ell \leq f$ and assume that $\sigma_{\mathcal{D}_b^\ell} \subseteq \sigma_\varphi$. To show that $\sigma^\ell \subseteq \sigma_\varphi$, recall that σ^ℓ is the concatenation of $\sigma_{\mathcal{D}_e^\ell}[0, i^\ell - 1]$ and $(\sigma_{\mathcal{D}_e^\ell}[i^\ell, j^\ell])^\omega$. As $\sigma_{\mathcal{D}_b^\ell} \subseteq \sigma_\varphi$ and \mathcal{D}_e^ℓ is obtained by applying (several times) `ApplyRules` to \mathcal{D}_b^ℓ , clearly $\sigma_{\mathcal{D}_e^\ell} \subseteq \sigma_\varphi$. Thus $\sigma_{\mathcal{D}_e^\ell}[0, i^\ell - 1]$ is contained in σ_φ . To show the inclusion for the second part of σ^ℓ , we will prove inductively on $n \in \mathbb{N}$ that $\sigma_{\mathcal{D}_e^\ell}, t \models p$ implies that $\sigma_\varphi, (t + n \cdot \text{per}) \models p$, for any $t \geq i^\ell$ (in particular for all $t \in [i^\ell, j^\ell]$), and $p \in \text{PROP}$, where `per` is the number of points in the period $[i^\ell, j^\ell]$. The base (for $n = 0$) holds as $\sigma_{\mathcal{D}_e^\ell} \subseteq \sigma_\varphi$. The inductive step for $t = i^\ell$ holds by the inductive assumption and saturation of \mathcal{D}_e^ℓ (implying that $\sigma_{\mathcal{D}_e^\ell}$ satisfies the same propositions at i^ℓ and $j^\ell + 1$). For $t > i^\ell$, we recall that \mathcal{D}_e^ℓ is obtained by applying (several times) `ApplyRules` to \mathcal{D}_b^ℓ , so we prove inductively on $m \in \mathbb{N}$ that $T_\Pi^m(\sigma_{\mathcal{D}_e^\ell}), t' \models p'$, implies that $\sigma_\varphi, (t' + n \cdot \text{per}) \models p'$, for all $t' > i^\ell$ and $p' \in \text{PROP}$. For the base ($m = 0$), since $t' > k_{\max}^\ell$, we get $\sigma_{\mathcal{D}_e^\ell}, t' \models Gp'$, and so $\sigma_{\mathcal{D}_e^\ell} \subseteq \sigma_\varphi$ implies that $\sigma_\varphi, (t' + n \cdot \text{per}) \models p'$. In the inductive step ($m > 0$) we assume without loss of generality that $T_\Pi^{m-1}(\sigma_{\mathcal{D}_e^\ell}), t' \not\models p'$, so $T_\Pi^m(\sigma_{\mathcal{D}_e^\ell}), t' \models p'$ by an application of a rule of one of the five types occurring in the normal-form Π . We can show that in each case the inductive assumption implies that $\sigma_\varphi, (t' + n \cdot \text{per}) \models p'$, as required.

Now, we consider Algorithm 2 running on φ (consisting of Π and \mathcal{D}) which can have \perp -rules and we let φ' be obtained by deleting all \perp -rules, so $\sigma_{\varphi'} = \mathcal{C}_{\Pi, \mathcal{D}}$.

If φ is satisfiable, then $\mathcal{C}_{\Pi, \mathcal{D}}$ satisfies all \perp -rules in Π . By the proof of Inclusion (\supseteq), $\mathcal{C}_{\Pi, \mathcal{D}} \models \mathcal{D}_b^\ell$ for each ℓ . Hence, no \perp -rule applies to any \mathcal{D}_b^ℓ , and the algorithm never returns `unsat` in Line ???. Thus the algorithm behaves as if it was running on φ' , and so, it returns the shortest representation of $\sigma_{\varphi'}$. Since φ is satisfiable, we get $\sigma_\varphi = \mathcal{C}_{\Pi, \mathcal{D}} = \sigma_{\varphi'}$, so the output is also the shortest representation of σ_φ .

If φ is unsatisfiable, we need to show that `unsat` is returned. Towards a contradiction suppose that some (u, w) is returned. Thus, the same (u, w) is returned if the algorithm runs on φ' , and so $uw^\omega = \sigma_{\varphi'} = \mathcal{C}_{\Pi, \mathcal{D}}$. Since φ is unsatisfiable, some of its \perp -rules is not satisfied in $\mathcal{C}_{\Pi, \mathcal{D}}$. Thus, in the last iteration the algorithm (running on φ) outputs `unsat` in Line ???. This, however, contradicts the assumption that (u, w) is returned. \square

We observe that $\text{LTL}_{\text{horn}}^{G,X}$, and even $\text{LTL}_{\text{horn}}^X$, allows us to simulate a binary counter (similarly to the proof of Proposition 8), so the exponential blow-up is unavoidable. In the next section we experimentally evaluate of our algorithms.

7 Experimental Evaluation

In this section, we will discuss the implementation of our algorithms and their experimental evaluation.

Implementation. We implemented both of our algorithms for computing the intersection of traces, namely




Figure 2: Evaluation of our automata-based Algorithm 1 and materialisation-based Algorithm 2; Figure 2(a) compares the two implementations running on Horn formulas; Figure 2(b) displays runtimes of the materialisation-based approach on Horn formulas; Figure 2(c) shows runtimes of the automata-based approach running on formulas from both the LTL (red crosses) and the Horn benchmarks (blue dots)

the automata-based Algorithm 1 for full LTL and the materialisation-based Algorithm 2 designed for $LTL_{horn}^{G,X}$. The automata-based implementation uses the state-of-the-art `spot` platform (Duret-Lutz et al. 2022) for constructing a Büchi automata corresponding to LTL-formulas (required in Line 1 of Algorithm 1). We wrote the implementations in Python and made them, together with all the benchmarks used in the experiments, available in open access.²

Benchmarks. Since computing the intersection of traces is a new task introduced in this paper, there were no benchmarks designed for this setting. Moreover, to the best of our knowledge, there are no benchmarks for formulas in the Horn forms we consider in the paper. Therefore, we randomly generated formulas required for testing our implementations. Our benchmarks consist of 100 LTL-formulas and another 100 $LTL_{horn}^{G,X}$ -formulas, as described below.

For LTL-formulas we used the random generator `randltl` provided in `spot` (Duret-Lutz et al. 2022). We set the parameter for the maximal number of propositions mentioned in a formula to 5 and the parameter ‘random tree size’, restricting the depth of the syntax tree of a formula, to values between 150 and 250. This allowed us to generate formulas of size varying from 1 to 461.

For $LTL_{horn}^{G,X}$ -formulas we had to implement a dedicated generator which provides pairs consisting of a dataset and a program. We generated such formulas randomly based on the following parameters: the number of propositions (ranging from 1 to 50), the number of facts in the dataset for each proposition (ranging from 1 to 12), the maximal nesting of X operator in the dataset (ranging from 0 to 22), and the number of rules in the program (ranging from 1 to 100). The obtained formulas have size varying from 28 to 13014.

Experiments. We ran our automata-based implementation on all of the benchmark formulas (so we used both LTL- and $LTL_{horn}^{G,X}$ -formulas), whereas the materialisation-based

implementation could be ran only on the $LTL_{G,X}^{horn}$ benchmark. We run the implementations 10 times for each formula and recorded the average runtime. Times exceeding 10 seconds were considered timeouts. All experiments were performed on a machine with an Intel® Core™ i7-8665U CPU@1.90Ghz and 16Gb RAM under Ubuntu Linux running inside a Windows 11 Subsystem for Linux.

The obtained results are presented in Figure 2. The scatter plot in Figure 2(a) shows that the materialisation-based implementation is significantly faster on every formula and always terminates within 0.1s, whereas the automata-based implementation times out on around half of the formulas (the time out often occurs already at the stage of constructing the automaton). Therefore, it is highly beneficial to use our dedicated materialisation-based implementation whenever the input formula falls into the Horn classes we considered. Figure 2(b) shows that the materialisation-based implementation scales well in practice and does not exhibit the exponential behaviour, which is in a sharp contrast with the results in Figure 2(c) for the automata-based approach.

8 Conclusions and Future Work

Motivated by applications in temporal knowledge representation, we introduced the problem of computing all facts entailed by an LTL-formula. We showed that the problem is doubly exponential, but becomes significantly easier in Horn fragments of LTL—in some cases it is tractable. We showed two optimal algorithms solving the problem for full LTL and its Horn fragments, respectively. We have implemented our algorithms and evaluated their performance. The results show a particularly good performance of the materialisation-based approach and suggest its practical usefulness.

In future, we plan to extend our materialisation-based approach to other classes of Horn formulas (e.g., allowing for past operators and the U operator in rule bodies). It would also be interesting to try to apply our approach to other temporal logics, for instance, to real time logics, metric temporal logics, and temporal description logics.

²The code and all benchmarks can be obtained in open access from <https://github.com/chalwidz/LTLIntersectionOfTraces>.

Acknowledgments

This research was supported in whole or in part by the EPSRC projects OASIS (EP/S032347/1), ConCuR (EP/V050869/1) and UK FIRES (EP/S019111/1), the SIR-IUS Centre for Scalable Data Access, and Samsung Research UK. The third author is supported by the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (Grant agreement No. 852769, ARiAT). For the purpose of Open Access, the authors have applied a CC BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission.

References

- Abadi, M., and Manna, Z. 1989. Temporal logic programming. *J. Symb. Comput.* 8(3):277–295.
- Aguado, F.; Cabalar, P.; Diéguez, M.; Pérez, G.; and Vidal, C. 2013. Temporal equilibrium logic: A survey. *J. Appl. Non-Class. Log.* 23(1-2):2–24.
- Alharby, M., and Van Moorsel, A. 2017. Blockchain-based smart contracts: A systematic mapping study. arXiv:1710.06372.
- Alur, R., and La Torre, S. 2004. Deterministic generators and games for LTL fragments. *ACM Trans. Comput. Log.* 5(1):1–25.
- Artale, A.; Kontchakov, R.; Ryzhikov, V.; and Zakharyashev, M. 2013. The complexity of clausal fragments of LTL. In *Proc. of LPAR*, 35–52.
- Artale, A.; Kontchakov, R.; Kovtunova, A.; Ryzhikov, V.; Wolter, F.; and Zakharyashev, M. 2021. First-order rewritability of ontology-mediated queries in linear temporal logic. *Artif. Intell.* 299:103536.
- Brandt, S.; Kalayci, E. G.; Kontchakov, R.; Ryzhikov, V.; Xiao, G.; and Zakharyashev, M. 2017. Ontology-based data access with a Horn fragment of metric temporal logic. In *Proc. of AAI*, 1070–1076.
- Bresolin, D.; Kurucz, A.; Muñoz-Velasco, E.; Ryzhikov, V.; Sciavicco, G.; and Zakharyashev, M. 2017. Horn fragments of the Halpern-Shoham interval temporal logic. *ACM Trans. Comput. Log.* 18(3):1–39.
- Chen, C.-C., and Lin, I.-P. 1993. The computational complexity of satisfiability of temporal Horn formulas in propositional linear-time temporal logic. *Inf. Process. Lett.* 45(3):131–136.
- Crochemore, M., and Rytter, W. 1994. *Text algorithms*. Oxford University Press.
- De Giacomo, G., and Vardi, M. Y. 2013. Linear temporal logic and linear dynamic logic on finite traces. In *Proc. of IJCAI*, 854–860.
- Demri, S., and Schnoebelen, P. 2002. The complexity of propositional linear temporal logics in simple cases. *Inf. Comput.* 174(1):84–103.
- Demri, S.; Goranko, V.; and Lange, M. 2016. *Temporal logics in computer science: Finite-state systems*. Cambridge University Press.
- Duret-Lutz, A.; Renault, E.; Colange, M.; Renkin, F.; Aisse, A. G.; Schlehuber-Caissier, P.; Medioni, T.; Martin, A.; Dubois, J.; Gillard, C.; and Lauko, H. 2022. From Spot 2.0 to Spot 2.10: What's new? In *Proc. of CAV*, 174–187.
- Fisher, M. 1991. A resolution method for temporal logic. In *Proc. of IJCAI*, volume 91, 99–104.
- Gabbay, D. 1987. Modal and temporal logic programming. In Galton, A., ed., *Temporal logics and their applications*. Academic Press. 197–237.
- Geatti, L.; Gigante, N.; and Montanari, A. 2021. Black: A fast, flexible and reliable LTL satisfiability checker. In *Proc. of OVERLAY*, volume 2987, 7–12.
- Lécué, F. 2017. Deep dive on smart cities by scaling reasoning and interpreting the semantics of IoT. In *Proc. of EGC*, 7–8.
- Li, J.; Yao, Y.; Pu, G.; Zhang, L.; and He, J. 2014. Aalta: an LTL satisfiability checker over infinite/finite traces. In *Proc. of SIGSOFT*, 731–734.
- Lichtenstein, O., and Pnueli, A. 2000. Propositional temporal logics: Decidability and completeness. *Log. J. IGPL* 8(1):55–85.
- Lutz, C.; Wolter, F.; and Zakharyashev, M. 2008. Temporal description logics: A survey. In *Proc. of TIME*, 3–14.
- Münz, G., and Carle, G. 2007. Real-time analysis of flow data for network attack detection. In *Proc. of IFIP/IEEE IM*, 100–108.
- Nair, M. 1982. On chebyshev-type inequalities for primes. *Am. Math. Mon.* 89(2):126–129.
- Nuti, G.; Mirghaemi, M.; Treleaven, P. C.; and Yingsaeree, C. 2011. Algorithmic trading. *IEEE Comput.* 44(11):61–69.
- Orgun, M. A., and Ma, W. 2005. An overview of temporal and modal logic programming. In *Proc. of ICTL*, 445–479.
- Pnueli, A. 1977. The temporal logic of programs. In *Proc. of FOCS*, 46–57.
- Ronca, A.; Kaminski, M.; Cuenca Grau, B.; Motik, B.; and Horrocks, I. 2018. Stream reasoning in temporal Datalog. In *Proc. of AAI*, 1941–1948.
- Ryzhikov, V.; Wałęga, P. A.; and Zakharyashev, M. 2020. Temporal ontology-mediated queries and first-order rewritability: A short course. In *Proc. of Reasoning Web. Declarative AI*, 109–148.
- Sistla, A. P., and Clarke, E. M. 1985. The complexity of propositional linear temporal logics. *J. ACM* 32(3):733–749.
- Vardi, M. Y., and Wolper, P. 1994. Reasoning about infinite computations. *Inf. Comput.* 115(1):1–37.
- Vardi, M. Y. 2005. An automata-theoretic approach to linear temporal logic. In Möller, F., and Birtwistle, G., eds., *Logics for concurrency: Structure versus automata*. Springer. 238–266.
- Vardi, M. Y. 2009. From philosophical to industrial logics. In *Proc. of ICLA*, 89–115.
- Wałęga, P. A. 2023. Computational complexity of hybrid interval temporal logics. *Ann. Pure Appl. Log.* 174(1):103165.

Wałęga, P. A.; Cuenca Grau, B.; Kaminski, M.; and Kostylev, E. V. 2019. DatalogMTL: Computational complexity and expressive power. In *Proc. of IJCAI*, 1886–1892.

Wałęga, P. A.; Cuenca Grau, B.; Kaminski, M.; and Kostylev, E. V. 2020. Tractable fragments of Datalog with metric temporal operators. In *Proc. of IJCAI*, 1919–1925.

Wolper, P. 1985. The tableau method for temporal logic: An overview. *Log. et Anal.* 119–136.

A Appendix

A.1 Proofs for Section 2

Proposition 2. *Each LTL-formula φ can be transformed in logarithmic space into a formula φ' in clausal form, such that φ and φ' entail the same LTL-formulas over the signature of φ .*

Proof. The transformation ‘renames’ subformulas ψ of φ to fresh propositions of the form a_ψ and introduces clauses which guarantee that each ψ and the corresponding a_ψ hold in the same time points. In particular, we let

$$\varphi' = a_\varphi \wedge \bigwedge_{\psi \text{ a subformula of } \varphi} \tau(\psi),$$

where the transformation τ is defined as follows, for any subformulas $p \in \text{PROP}$, ψ, ψ_1, ψ_2 of φ :

$$\begin{aligned} \tau(p) &= \text{G}(a_p \leftrightarrow p), \\ \tau(\neg\psi) &= \text{G}(a_{\neg\psi} \leftrightarrow \neg a_\psi), \\ \tau(\psi_1 \vee \psi_2) &= \text{G}(a_{\psi_1 \vee \psi_2} \leftrightarrow a_{\psi_1} \vee a_{\psi_2}), \\ \tau(\mathbf{X}\psi) &= \text{G}(a_{\mathbf{X}\psi} \leftrightarrow \mathbf{X}a_\psi), \\ \tau(\psi_1 \mathbf{U}\psi_2) &= \text{G}(a_{\psi_1 \mathbf{U}\psi_2} \leftrightarrow a_{\psi_1} \mathbf{U}a_{\psi_2}), \end{aligned}$$

where $\text{G}(A \leftrightarrow B)$ is an abbreviation for $\text{G}(\neg A \vee B) \wedge \text{G}(\neg B \vee A)$, whereas $\text{G}(A \leftrightarrow \neg B)$ is an abbreviation for $\text{G}(\neg A \vee \neg B) \wedge \text{G}(B \vee A)$, and $\text{G}(A \leftrightarrow B \vee C)$ is an abbreviation for $\text{G}(\neg A \vee B \vee C) \wedge \text{G}(\neg B \vee A) \wedge \text{G}(\neg C \vee A)$.

To finish the proof, we show that for every LTL-formula χ over the signature of φ it holds that φ entails χ if and only if φ' entails χ .

(\Rightarrow) First assume that φ entails χ . Fix a trace σ such that $\sigma, 0 \models \varphi'$ (note that if there is no such σ , then φ' entails χ trivially). By the form of φ' , we have $\sigma, 0 \models a_\varphi$. Thus, by the form of the clauses in φ' , we obtain that $\sigma, 0 \models \varphi$. Hence, by the assumption, $\sigma, 0 \models \chi$, which shows that φ' entails χ .

(\Leftarrow) Now assume that φ' entails χ and let us fix σ such that $\sigma, 0 \models \varphi$. We let σ' be similar to σ except that we set the valuation of propositions a_ψ , for all subformulas ψ of φ , such that for each $i \in \mathbb{N}$ we have $\sigma', i \models a_\psi$ if and only if $\sigma, i \models \psi$. Hence, by the form of φ' , we obtain that $\sigma', 0 \models \varphi'$ and so, by the assumption that φ' entails χ , we have that $\sigma', 0 \models \chi$. Finally, as χ is over the signature of φ , it follows that χ does not mention propositions of the form a_ψ , and so $\sigma', 0 \models \chi$ implies that $\sigma, 0 \models \chi$. Therefore, φ entails χ . \square

A.2 Proofs for Section 3

Proposition 4. *For any LTL-formula φ the intersection of all its traces, σ_φ , is of the form uw^ω , for some finite traces u and w .*

Proof. To finish the proof sketched in the main body of the paper, we need to define the formula $\text{tr}_\varphi(X_1, \dots, X_n)$, whose aim is to state that φ is satisfied at 0 in the trace in which p_i hold in positions from X_i .

To this end, we let $\text{tr}_\varphi(X_1, \dots, X_n) = \tau_\varphi(X_1, \dots, X_n, 0)$ and we define $\tau_\psi(X_1, \dots, X_n, x)$, for all LTL-formulas ψ , inductively on the structure of the formula. In particular, for any proposition p_i and LTL-formulas μ and μ' we let:

$$\begin{aligned} \tau_{p_i}(X_1, \dots, X_n, x) &= x \in X_i, \\ \tau_{\neg\mu}(X_1, \dots, X_n, x) &= \neg\tau_\mu(X_1, \dots, X_n, x), \\ \tau_{\mu \vee \mu'}(X_1, \dots, X_n, x) &= \tau_\mu(X_1, \dots, X_n, x) \vee \tau_{\mu'}(X_1, \dots, X_n, x), \\ \tau_{\mathbf{X}\mu}(X_1, \dots, X_n, x) &= \tau_\mu(X_1, \dots, X_n, x + 1), \\ \tau_{\text{G}\mu}(X_1, \dots, X_n, x) &= \forall y (x \leq y \rightarrow \tau_\mu(X_1, \dots, X_n, y)). \end{aligned}$$

Now, for any subsets P_1, \dots, P_n of \mathbb{N} , we let σ_{P_1, \dots, P_n} be the trace in which each proposition p_i is satisfied in the positions belonging to P_i . Then, by the definition of τ , we obtain that for each LTL-formula ψ and $i \in \mathbb{N}$ we have:

$$\tau_\psi(P_1, \dots, P_n, i) \text{ holds true in SIS} \quad \text{iff} \quad \sigma_{P_1, \dots, P_n}, i \models \psi.$$

In particular, $\tau_\varphi(P_1, \dots, P_n, 0)$ holds true in SIS if and only if $\sigma_{P_1, \dots, P_n}, 0 \models \varphi$, as required. \square

In Algorithm 3 we refer to an algorithm computing the shortest representation (u', w') of uw^ω , which works in linear time.

Algorithm 3: Computing the shortest representation

Input: words u and w **Output:** the shortest representation (u', w') of uw^ω

```
1  $i := 0$ ;  
2 repeat  
3    $i := i + 1$ ; // compute the shortest length  $i$  of period  
4    $a := \text{lcm}(i, |w|)/i$ ;  
5    $b := \text{lcm}(i, |w|)/|w|$ ;  
6 until  $(w[0, i - 1])^a = w^b$ ;  
7  $j := |u|$ ;  
8 repeat  
9    $j := j - 1$ ; // compute the last position  $j$  of the prefix  
10 until  $uw(j) \neq uw(j + i)$  or  $j = 0$ ;  
11  $u' := u[0, j]$ ;  
12  $w' := uw[j + 1, j + i]$ ;  
13 return  $(u', w')$ ;
```

A.3 Proofs for Section 5

Theorem 9. For every $\text{LTL}_{\text{horn}}^{\text{G}}$ -formula φ , composed of a program Π and dataset \mathcal{D} , we have that $\mathfrak{C}_{\Pi, \mathcal{D}} = T_{\Pi}^{|\varphi|^2}(\sigma_{\mathcal{D}})$.

Proof. To complete the proof sketch from the main body of the paper it remains to show by a transfinite induction on ordinals α that (i) $T_{\Pi}^{\alpha}(\sigma_{\mathcal{D}})$ satisfies the same propositions in all time points belonging to the same segment and that (ii) $T_{\Pi}^{\alpha}(\sigma_{\mathcal{D}})$ does not satisfy any proposition in the segment $[0, k_1)$.

The inductive base (for $\alpha = 0$) holds directly by the definition of segments. Next, we show the inductive step for a successor ordinal α . By the definition, we have that $T_{\Pi}^{\alpha}(\sigma_{\mathcal{D}}) = T_{\Pi}(T_{\Pi}^{\alpha-1}(\sigma_{\mathcal{D}}))$. As Π is an $\text{LTL}_{\text{horn}}^{\text{G}}$ -program, all its body atoms are of the form p or $\text{G}p$, for $p \in \text{PROP}$. Hence, by the inductive assumption, if for some rule $r \in \Pi$ and a time point i we have $T_{\Pi}^{\alpha-1}(\sigma_{\mathcal{D}}), i \models \text{body}(r)$, then $T_{\Pi}^{\alpha-1}(\sigma_{\mathcal{D}}), j \models \text{body}(r)$ for all the time points j belonging to the same segment as i and, moreover, $i \notin [0, k_1)$. Furthermore, as Π is an $\text{LTL}_{\text{horn}}^{\text{G}}$ -program the head of r is either of the form p or $\text{G}p$, for some $p \in \text{PROP}$. Thus $T_{\Pi}^{\alpha}(\sigma_{\mathcal{D}})$, which is the minimal trace which satisfies this head at i and contains $T_{\Pi}^{\alpha-1}(\sigma_{\mathcal{D}})$, needs to satisfy both Statements (i) and (ii). Finally, we observe that the inductive step for a limit ordinal α holds trivially since, by the definition, $T_{\Pi}^{\alpha}(\sigma_{\mathcal{D}}) = \bigcup_{\beta < \alpha} T_{\Pi}^{\beta}(\sigma_{\mathcal{D}})$. \square

Theorem 12. For every $\text{LTL}_{\text{horn}}^{\text{G}, \text{X}}$ program Π and dataset \mathcal{D} , we have $\mathfrak{C}_{\Pi, \mathcal{D}} = T_{\Pi}^{\omega \cdot (k+1)}(\sigma_{\mathcal{D}})$, where k is the number of rules in Π which mention G in their bodies.

Proof. To complete the proof sketch in the body of the paper, it remains to show that there exists a unique minimal (finite) dataset \mathcal{D}_r such that $\sigma_{\mathcal{D}_r} \subseteq T_{\Pi}^{\alpha}(\sigma_{\mathcal{D}})$ and $\sigma_{\mathcal{D}_r} \cup T_{\Pi}^{\beta}(\sigma_{\mathcal{D}}), i \models \text{body}(r)$.

To construct \mathcal{D}_r , we observe that each body atom of the $\text{LTL}_{\text{horn}}^{\text{G}, \text{X}}$ rule r can be written as $\text{X}^k p$ or $\text{X}^k \text{G}p$, for some $k \in \mathbb{N}$ and $p \in \text{PROP}$. For every such body atom λ we will construct a minimal dataset \mathcal{D}_{λ} which needs to be added to $T_{\Pi}^{\beta}(\sigma_{\mathcal{D}})$ so that λ will be satisfied at i . If already $T_{\Pi}^{\beta}(\sigma_{\mathcal{D}}), i \models \lambda$, we do not need to add any facts, so we let $\mathcal{D}_{\lambda} = \emptyset$. Otherwise, we consider the cases when λ is of the form $\text{X}^k p$ and $\text{X}^k \text{G}p$, separately. If $\lambda = \text{X}^k p$, we let $\mathcal{D}_{\lambda} = \{\text{X}^k p\}$. If $\lambda = \text{X}^k \text{G}p$, the definition of \mathcal{D}_{λ} is more complex. We recall that $T_{\Pi}^{\beta}(\sigma_{\mathcal{D}}), j \models \text{body}(r)$, for $i \neq j$, so $T_{\Pi}^{\beta}(\sigma_{\mathcal{D}}), j \models \text{X}^k \text{G}p$. Thus, if $j < i$ we have $T_{\Pi}^{\beta}(\sigma_{\mathcal{D}}), i \models \text{X}^k \text{G}p$, and so, $\mathcal{D}_{\lambda} = \emptyset$. Otherwise $i < j$ and to make $\text{X}^k \text{G}p$ satisfied at i we need to add to $T_{\Pi}^{\beta}(\sigma_{\mathcal{D}})$ the facts $\mathcal{D}_{\lambda} = \{\text{X}^{\ell} p \mid i + k \leq \ell < j + k\}$.

Finally, we let \mathcal{D}_r be the union of \mathcal{D}_{λ} for all body atoms λ in r . By the definition of \mathcal{D}_r we obtain that $\sigma_{\mathcal{D}_r} \subseteq T_{\Pi}^{\alpha}(\sigma_{\mathcal{D}})$ and $\sigma_{\mathcal{D}_r} \cup T_{\Pi}^{\beta}(\sigma_{\mathcal{D}}), i \models \text{body}(r)$, and moreover, \mathcal{D}_r is the minimal set which satisfies these properties. \square

A.4 Proofs for Section 6

Without loss of generality let's assume that φ is a $\text{LTL}_{\text{horn}}^{\text{G}, \text{X}}$ -formula such that Π is the program and \mathcal{D} is the dataset in φ . Below we show how to construct a program $\text{Norm}(\Pi)$ in normal from Π .

1. Initialise $\text{Norm}(\Pi)$ as an empty program.
2. For each rule $\lambda \leftarrow \lambda_1 \wedge \dots \wedge \lambda_n$ in Π such that λ_i does not contain \perp , for each $i \in \{1, \dots, n\}$, add an auxiliary rule $p^{\lambda} \leftarrow p^{\lambda_1} \wedge \dots \wedge p^{\lambda_n}$ to an auxiliary set Π' , where $p^{\lambda}, p^{\lambda_1}, \dots, p^{\lambda_n}$ are fresh propositional atoms not occurring in Π .

3. For each atom $\lambda = \otimes_n \cdots \otimes_1 q$, occurring in the body of a rule in Π' , where $\otimes_1, \dots, \otimes_n \in \{G, X\}$ and $q \in \text{PROP}$, add the following rules to $\text{Norm}(\Pi)$:

- $p^{\otimes_1 q} \leftarrow \otimes_1 q$ and
- $p^{\otimes_{k+1} \cdots \otimes_1 q} \leftarrow \otimes_{k+1} p^{\otimes_k \cdots \otimes_1 q}$, $1 \leq k < n$,

where $p^{\otimes_i \cdots \otimes_1 q}$, for each $k \in \{1, \dots, n\}$, is a fresh propositional atom not occurring in Π .

4. For each atom of the form $\otimes_n \cdots \otimes_1 q$, occurring in the head of a rule in Π' , where $\otimes_1, \dots, \otimes_n \in \{G, X\}$ and $q \in \text{PROP} \cup \{\perp\}$ add the following rules to $\text{Norm}(\Pi)$:

- if $q \in \text{PROP}$:
 - $\otimes_1 q \leftarrow p^{\otimes_1 q}$ and
 - $\otimes_{k+1} p^{\otimes_k \cdots \otimes_1 q} \leftarrow p^{\otimes_{k+1} \cdots \otimes_1 q}$, $1 \leq k < n$,
- if $q = \perp$:
 - $\perp \leftarrow p^\perp \wedge p^{\perp}$ and
 - $\otimes_1 p^\perp \leftarrow p^{\otimes_1 \perp}$, and
 - $\otimes_{k+1} p^{\otimes_k \cdots \otimes_1 \perp} \leftarrow p^{\otimes_{k+1} \cdots \otimes_1 \perp}$, $1 \leq k < n$,

where p^\perp and $p^{\otimes_k \cdots \otimes_1 q}$, for each $k \in \{1, \dots, n\}$, is a fresh propositional atom not occurring in Π .

5. For each rule $r = q \leftarrow p_1 \wedge \dots \wedge p_n$ in Π or in Π' add the following rules to $\text{Norm}(\Pi)$:

- $q \leftarrow p^{p_1} \wedge p^{p_1}$, if $n = 1$;
- $q \leftarrow p^{p_1 \wedge \dots \wedge p_{n-1}} \wedge p^{p_n}$ and $p^{p_1 \wedge \dots \wedge p_{k+1}} \leftarrow p^{p_1 \wedge \dots \wedge p_k} \wedge p^{p_{k+1}}$, $1 \leq k < n - 1$, if $n > 1$,

where, p^φ is a fresh propositional atom not occurring in Π .

Theorem 17. Let φ be an $LTL_{horn}^{G,X}$ -formula consisting of a program Π and a dataset \mathcal{D} and let φ' be φ with Π replaced with $\text{Norm}(\Pi)$. Then $\sigma_{\varphi'}$ is a conservative extension of σ_φ .

Proof. Let $\varphi, \varphi', \Pi, \text{Norm}(\Pi)$, and \mathcal{D} be as in the theorem. First, we will show by induction on subsequent applications of the operator T_Π to $\sigma_\mathcal{D}$ that, if $T_\Pi^k(\sigma_\mathcal{D}), i \models p$, for $k \in \mathbb{N}$ and $p \in \text{PROP}$, then there exists $\ell \in \mathbb{N}$ such that $T_{\text{Norm}(\Pi)}^\ell(\sigma_\mathcal{D}), i \models p$. For a trace σ , interval ϱ , and $p \in \text{PROP}$, by $\sigma, \varrho \models p$ we mean $\sigma, i' \models p$, for all $i' \in \varrho$.

Base case: $k = 0$ We have $T_\Pi^0(\sigma_\mathcal{D}) = \sigma_\mathcal{D} = T_{\text{Norm}(\Pi)}^0(\sigma_\mathcal{D})$, and so, the claim holds.

Inductive step: > 0 Assume that the claim holds for all $k' < k$. We will show that it also holds for k . Let $T_\Pi^k(\sigma_\mathcal{D}), i \models p$ for $p \in \text{PROP}$. Without loss of generality let's assume that $T_\Pi^k(\sigma_\mathcal{D}), i \models p$, yet $T_\Pi^{k-1}(\sigma_\mathcal{D}) \not\models p$. It means that there is a rule $r = \lambda \leftarrow \lambda_1 \wedge \dots \wedge \lambda_n$ in Π , propositional atoms p_1, \dots, p_n , and time intervals $\varrho, \varrho_1, \dots, \varrho_n$ such that $T_\Pi^{k-1}(\sigma_\mathcal{D}), \varrho_j \models p_j$, for $1 \leq j \leq n$, $T_\Pi^{k-1}(\sigma_\mathcal{D}), \varrho_j \models p_j$ entails $T_\Pi^{k-1}(\sigma_\mathcal{D}), i' \models \lambda_j$, for $1 \leq j \leq n$, and $T_\Pi^{k-1}(\sigma_\mathcal{D}), i' \models \lambda$ entails $T_\Pi^{k-1}(\sigma_\mathcal{D}), i \models p$. Notice that, for any temporal atom λ , if λ contains \perp , then λ is unsatisfiable. Therefore, it must indeed be the case that, for each $i \in \{1, \dots, n\}$, λ_i does not contain \perp . By the inductive hypothesis, there exist ℓ such that $T_{\text{Norm}(\Pi)}^\ell(\sigma_\mathcal{D}), \varrho_j \models p_j$, for $1 \leq j \leq n$, and so. By the construction of $\text{Norm}(\Pi)$, $T_{\text{Norm}(\Pi)}^{\ell + \text{depth}(\Pi)}(\sigma_\mathcal{D}), i' \models p^{\lambda_j}$, for $1 \leq j \leq n$, where p^λ is simply λ if λ is a propositional atom. Then, again by the construction of $\text{Norm}(\Pi)$, $T_{\text{Norm}(\Pi)}^{\ell + 2\text{depth}(\Pi) + \text{width}(\Pi)}(\sigma_\mathcal{D}), i \models p$, which proves the claim of the inductive step.

Now let us show that, for any propositional atom p occurring in φ and any time point i , if $\sigma_{\varphi'}, i \models p$, then $\sigma_\varphi, i \models p$. We show, by induction on subsequent applications of the operator $T_{\text{Norm}(\Pi)}$ to $\sigma_\mathcal{D}$, that if $T_{\text{Norm}(\Pi)}^k(\sigma_\mathcal{D}), i \models p$, yet $T_{\text{Norm}(\Pi)}^{k-1}(\sigma_\mathcal{D}), i \not\models p$, then there exists $\ell \in \mathbb{N}$ such that $T_\Pi^\ell(\sigma_\mathcal{D}), i \models p$.

Base case: $k = 0$ We have $T_{\text{Norm}(\Pi)}^0(\sigma_\mathcal{D}) = \sigma_\mathcal{D} = T_\Pi^0(\sigma_\mathcal{D})$, and so, the claim holds.

Inductive step: $k > 0$ Assume that the claim holds for all $k' < k$. We will show that it also holds for k . Let $T_{\text{Norm}(\Pi)}^k(\sigma_\mathcal{D}), i \models p$, yet $T_{\text{Norm}(\Pi)}^{k-1}(\sigma_\mathcal{D}), i \not\models p$. By the construction of $\text{Norm}(\Pi)$, the only way of introducing p to $T_{\text{Norm}(\Pi)}^k(\sigma_\mathcal{D})(i)$, where p is a propositional atom occurring in Π , is by applying a rule r from point (3) or (5) of the normalisation algorithm which stems from $r' = \lambda \leftarrow \lambda_1 \wedge \dots \wedge \lambda_n \in \Pi$. It means that there exists $\ell < k$ such that $T_{\text{Norm}(\Pi)}^\ell(\sigma_\mathcal{D}), \varrho_j \models p_j$, for $1 \leq j \leq n$, where p_1, \dots, p_n are propositional atoms from Π occurring in r and all the rules stemming from r' that trigger the above-mentioned rules resulting in p being added to $T_{\text{Norm}(\Pi)}^k(\sigma_\mathcal{D})(i)$. By the inductive hypothesis there exists ℓ such that $T_\Pi^\ell(\sigma_\mathcal{D}), \varrho_j \models p_j$, for $1 \leq j \leq n$. Since $T_\Pi^\ell(\sigma_\mathcal{D}), \varrho_j \models p_j$ entails $T_\Pi^\ell(\sigma_\mathcal{D}), i' \models \lambda_j$, for $1 \leq j \leq n$, and a rule $\lambda \leftarrow \lambda_1 \wedge \dots \wedge \lambda_n$ is in Π , we have $T_\Pi^\ell(\sigma_\mathcal{D}), i' \models \lambda$. Since $T_\Pi^\ell(\sigma_\mathcal{D}), i' \models \lambda$ entails $T_\Pi^\ell(\sigma_\mathcal{D}), i \models p$, we obtain the conclusion. \square

Theorem 15. *The output (u, w) of Algorithm 2, running on a formula φ , is the shortest representation of σ_φ .*

Proof. To finish the proof of the inclusion $\sigma_\varphi \subseteq \sigma^f$ it remains to show that Statements (i) and (ii) hold for $\lambda = Gp$. We start by showing the left-to-right implications of Statements (i) and (ii). If $\sigma^f, t \models Gp$, then $\sigma^f, s \models p$ for all $s \in [t, j^f]$ and for all $s \in [j^f+1, j^f+1+(j^f-i^f)]$. By the construction of σ^f , it implies that $\sigma_{\mathcal{D}_e^f}, s \models p$ for all $s \in [t, j^f]$ and for all $s \in [i^f, j^f]$, respectively. The latter implies that if $\sigma_{\mathcal{D}_e^f}, i^f \not\models Gp$, then $X^{i^f}Gp \in \mathcal{D}_\infty^f$. However, $\mathcal{D}_\infty^f = \emptyset$ (see Line 10), so $\sigma_{\mathcal{D}_e^f}, i^f \models Gp$. Thus $\sigma_{\mathcal{D}_e^f}, t \models Gp$. The right-to-left implication for Statement (i) holds by the definition of σ^f and for Statement (ii), $\sigma_{\mathcal{D}_e^f}, t' \models Gp$ implies that $\sigma_{\mathcal{D}_e^f}, j^f+1 \models Gp$ (as $t' \leq j^f$), so $\sigma_{\mathcal{D}_e^f}, i^f \models Gp$ (by saturation). Thus $\sigma^f, i^f \models Gp$ and $\sigma^f, t \models Gp$ (as $t \geq i^f$).

To finish the proof of the other inclusion, namely $\sigma^f \subseteq \sigma_\varphi$, we will show, by induction on $n \in \mathbb{N}$, that the following implication holds for any time point $t \geq i^\ell$ and $p \in \text{PROP}$:

$$\text{if } \sigma_{\mathcal{D}_e^\ell}, t \models p, \text{ then } \sigma_\varphi, t+n \cdot \text{per} \models p. \quad (\dagger)$$

Base case, $n = 0$:

Since $t+0 \cdot \text{per} = t$, (\dagger) follows from the inclusion $\sigma_{\mathcal{D}_e^\ell} \subseteq \sigma_\varphi$.

Inductive step, $n > 0$: We assume that (\dagger) holds for all $n' < n$. We will show that it holds also for n . Two cases may occur: $t = i^\ell$ and $t > i^\ell$.

If $t = i^\ell$, we observe that, by the assumption on the saturation of \mathcal{D}_e^ℓ , $\sigma_{\mathcal{D}_e^\ell}, j^\ell + 1 \models p$. By the inductive hypothesis we get that $\sigma_\varphi, j^\ell+1+(n-1) \cdot \text{per} \models p$. But $j^\ell+1$ is precisely $i^\ell + \text{per}$, and so, we obtain $\sigma_\varphi, t+n \cdot \text{per} \models p$.

Now, since \mathcal{D}_e^ℓ was obtained by m^ℓ applications of ApplyRules to \mathcal{D}_b^ℓ , we will show, by induction on $m \leq m^\ell$, that, for any $t > i^\ell$ and any $p \in \text{PROP}$:

$$\text{if } \sigma_{\text{ApplyRules}^m(\Pi, \mathcal{D}_b^\ell)}, t \models p, \text{ then } \sigma_\varphi, t+n \cdot \text{per} \models p. \quad (\ddagger)$$

Base case, $m = 0$: Let $p \in \text{PROP}$, $t > i^\ell$ and $\sigma_{\text{ApplyRules}^0(\Pi, \mathcal{D}_b^\ell)}, t \models p$. Since $\sigma_{\text{ApplyRules}^0(\Pi, \mathcal{D}_b^\ell)} = \sigma_{\mathcal{D}_b^\ell}$ and $t > i^\ell > k_{\max}^\ell$, where k_{\max}^ℓ is the maximal k such that an expression X^k occurs in \mathcal{D}_b^ℓ , we know that $\sigma_{\mathcal{D}_b^\ell}, t \models p$ if and only if $\sigma_{\mathcal{D}_b^\ell}, t \models Gp$. By the semantics of G we obtain $\sigma_{\mathcal{D}_b^\ell}, t+n \cdot \text{per} \models p$. By the assumption, $\sigma_{\mathcal{D}_b^\ell} \subseteq \sigma_\varphi$, whence the conclusion follows.

Inductive step, $m > 0$: Let $p \in \text{PROP}$, $t > i^\ell$ and $\sigma_{\text{ApplyRules}^m(\Pi, \mathcal{D}_b^\ell)}, t \models p$. Without loss of generality we may assume that $\sigma_{\text{ApplyRules}^{m-1}(\Pi, \mathcal{D}_b^\ell)}, t \not\models p$ and let's assume that for any propositional atom q , time point $t' > i^\ell$, and $m' < m$ the claim holds. We know that p was added to $\sigma_{\text{ApplyRules}^m(\Pi, \mathcal{D}_b^\ell)}(t)$ in an m th application of ApplyRules to \mathcal{D}_b^ℓ by firing a rule r . We consider the following cases:

$r = p \leftarrow q \wedge r$ If r was applied in an m th application of ApplyRules to \mathcal{D}_b^ℓ , it means that $\sigma_{\text{ApplyRules}^{m-1}(\Pi, \mathcal{D}_b^\ell)}, t \models q, r$.

By the inductive hypothesis we obtain $\sigma_\varphi, t+n \cdot \text{per} \models q, r$. Since σ_φ satisfies all the rules from Π , it follows that $\sigma_\varphi, t+n \cdot \text{per} \models p$.

$r = Xp \leftarrow q$ If $t = i^\ell+1$, it means that $\sigma_{\text{ApplyRules}^{m-1}(\Pi, \mathcal{D}_b^\ell)}, i^\ell \models q$, and so, $\sigma_{\mathcal{D}_e^\ell}, i^\ell \models q$. From the case for $t = i^\ell$, we know that it implies $\sigma_\varphi, i^\ell+n \cdot \text{per} \models q$. As σ_φ satisfies all the rules from Π , we obtain $\sigma_\varphi, i^\ell+n \cdot \text{per} \models Xp$, and so, $\sigma_\varphi, i^\ell+1+n \cdot \text{per} \models p$. Otherwise let's assume that $t > i^\ell+1$. Then, it must be the case that $\sigma_{\text{ApplyRules}^{m-1}(\Pi, \mathcal{D}_b^\ell)}, t-1 \models q$. Since $t-1 > i^\ell$, we can apply the inductive hypothesis, thus obtaining $\sigma_\varphi, t-1+n \cdot \text{per} \models q$. For the same reason as above, we finally get $\sigma_\varphi, t+n \cdot \text{per} \models p$.

$r = p \leftarrow Xq$ Since r was applied resulting in $\sigma_{\text{ApplyRules}^m(\Pi, \mathcal{D}_b^\ell)}, t \models p$, it must have been the case that $\sigma_{\text{ApplyRules}^{m-1}(\Pi, \mathcal{D}_b^\ell)}, t+1 \models q$. Of course, $t+1 > i^\ell$, so we can apply the inductive hypothesis, getting $\sigma_\varphi, t+1+n \cdot \text{per} \models q$. Since σ_φ satisfies all the rules from Π , we obtain $\sigma_\varphi, t+n \cdot \text{per} \models p$.

$r = Gp \leftarrow q$ Since $\sigma_{\text{ApplyRules}^m(\Pi, \mathcal{D}_b^\ell)}, t \models p$, but $\sigma_{\text{ApplyRules}^{m-1}(\Pi, \mathcal{D}_b^\ell)}, t \not\models p$, it must have been the case that $\sigma_{\text{ApplyRules}^{m-1}(\Pi, \mathcal{D}_b^\ell)}, t' \models q$, for some $t' \leq t$. By the semantics of G, we obtain $\sigma_{\text{ApplyRules}^m(\Pi, \mathcal{D}_b^\ell)}, t'' \models p$, for any $t'' \geq t'$, and so, $\sigma_{\text{ApplyRules}^m(\Pi, \mathcal{D}_b^\ell)}, t+n \cdot \text{per} \models p$, as $t+n \cdot \text{per} > t'$. Since $\sigma_{\text{ApplyRules}^m(\Pi, \mathcal{D}_b^\ell)} \subseteq \sigma_\varphi$, the claim follows.

$r = p \leftarrow Gq$ It must have been the case that $\sigma_{\text{ApplyRules}^{m-1}(\Pi, \mathcal{D}_b^\ell)}, t \models Gq$. By the semantics of G, $\sigma_{\text{ApplyRules}^{m-1}(\Pi, \mathcal{D}_b^\ell)}, t+n \cdot \text{per} \models Gq$, and so, $\sigma_{\text{ApplyRules}^m(\Pi, \mathcal{D}_b^\ell)}, t+n \cdot \text{per} \models p$. Since $\sigma_{\text{ApplyRules}^m(\Pi, \mathcal{D}_b^\ell)} \subseteq \sigma_\varphi$, we obtain the conclusion.

Since $\text{ApplyRules}^{m^\ell}(\Pi, \mathcal{D}_b^\ell) = \mathcal{D}_e^\ell$, we get the conclusion. \square

Proposition 18. *There is a family of $LTL_{\text{hor}n}^X$ -formulas φ_n , with $n \in \mathbb{N}$, such that each φ_n is of polynomial size in n and σ_{φ_n} has the shortest representation (u, w) with $|w| = 2^{\Omega(n)}$.*

Proof. We will use φ_n to encode a binary counter. Formula φ_n will use propositions $0, 1, \#, \triangleleft$ and the least traces σ_{φ_n} satisfying φ_n will be of the form

$$(\text{bin}_n(0) \# \text{bin}_n(1) \# \cdots \# \text{bin}_n(2^n - 1) \triangleleft)^\omega,$$

where $\text{bin}_n(m)$ is the encoding of m using n bits $b_0 b_1 \cdots b_{n-1} \in \{0, 1\}^n$ with the least-significant bit first. Thus the shortest representation (u, w) of σ_{φ_n} is such that $u = \emptyset$ and $|w| = (n + 1) \cdot 2^n$.

The definition of φ_n uses a similar idea to the one from the proof of Proposition 8. In particular, we let φ_n consist of a dataset \mathcal{D} and a program Π , where \mathcal{D} is the set containing facts $X^i 0$ for all $0 \leq i \leq n - 1$ and the fact $X^{n+1} \#$, whereas Π consists of the following rules using propositions $\text{Inc}_{n,i}$ and $\text{Equal}_{n,i}$, for all $0 \leq i \leq n$:

$$\begin{aligned} X^n \# &\rightarrow \text{Inc}_{n,0}, \\ \text{Inc}_{n,i} \wedge X^i 0 &\rightarrow X^{i+n+1} 1, \\ \text{Inc}_{n,i} \wedge X^i 0 &\rightarrow \text{Equal}_{n,i+1}, \\ \text{Inc}_{n,i} \wedge X^i 1 &\rightarrow X^{i+n+1} 0, \\ \text{Inc}_{n,i} \wedge X^i 1 &\rightarrow \text{Inc}_{n,i+1}, \\ \text{Equal}_{n,i} \wedge X^j 0 &\rightarrow X^{j+n+1} 0, & \text{for all } i \leq j \leq n - 1 \\ \text{Equal}_{n,i} \wedge X^j 1 &\rightarrow X^{j+n+1} 1, & \text{for all } i \leq j \leq n - 1 \\ X^1 1 \wedge X^2 1 \wedge \cdots \wedge X^n 1 &\rightarrow X^{n+1} \triangleleft, \\ \triangleleft &\rightarrow X^j 0, & \text{for all } 1 \leq j \leq n. \end{aligned}$$

To finish, we observe that the size of φ_n is polynomial in the size of n , as required. □