On Finite Domains in First-Order Linear Temporal Logic

Denis Kuperberg, Julien Brunel, David Chemouil

► To cite this version:

Denis Kuperberg, Julien Brunel, David Chemouil. On Finite Domains in First-Order Linear Temporal Logic. 14th International Symposium on Automated Technology for Verification and Analysis, Oct 2016, Chiba, Japan. 10.1007/978-3-319-46520-3_14 . hal-01343197

HAL Id: hal-01343197

<https://hal.science/hal-01343197>

Submitted on 7 Jul 2016

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

On Finite Domains in First-Order Linear Temporal Logic^{*}

Denis Kuperberg¹, Julien Brunel², and David Chemouil²

¹ TU Munich, Germany

² DTIM, Université fédérale de Toulouse, ONERA, France

Abstract. We consider First-Order Linear Temporal Logic (FO-LTL) over linear time. Inspired by the success of formal approaches based upon finite-model finders, such as Alloy, we focus on finding models with finite first-order domains for FO-LTL formulas, while retaining an infinite time domain. More precisely, we investigate the complexity of the following problem: given a formula φ and an integer n , is there a model of φ with domain of cardinality at most n ? We show that depending on the logic considered (FO or FO-LTL) and on the precise encoding of the problem, the problem is either NP-complete, NEXPTIME-complete, PSPACE-complete or EXPSpace-complete. In a second part, we exhibit cases where the Finite Model Property can be lifted from fragments of FO to their FO-LTL extension.

Keywords: FO, LTL, Finite model property, Bounded satisfiability, Fragments

Extended version of a paper presented at ATVA 2016. The final publication is available at link.springer.com.

1 Introduction

1.1 Context

First-Order Logic (FO) has proven to be useful in numerous applications in computer science such as formal specification, databases, ontology languages, etc. It is particularly well-suited to reason about objects of a domain, their relations and the properties they satisfy. However, since “full” FO is undecidable, the formal *verification* of properties implies a relaxation of the problem *e.g.* considering less expressive fragments. Thus, one can restrict the specification language (*e.g.* Prolog) or impose some form of interaction for verification (*e.g.* theorem provers, proof assistants).

Another form of trade-off is to keep the whole logic and full automation but to rely on a sound but incomplete decision procedure. For instance, the Alloy

^{*} Research partly funded by ANR/DGA project Cx (ref. ANR-13-ASTR-0006) and by *fondation STAE* project BRIefcaSE.

Analyzer³ for the Alloy [Jac06] language (based upon relational first-order logic) implements a *bounded-satisfiability* decision procedure. That is, the tool relies on a *finite-model finder*: it first bounds the number of objects in the domain and then runs a classical propositional SAT procedure [TJ07]. Thanks to the performance of modern SAT engines, this approach has shown to be very efficient in practice to find counterexamples quickly when assessing specifications. This is one of the reasons for the success of Alloy, in the formal methods community [Zav12, NRZ⁺15, BKMJ15].

However, in most software and systems specifications, one needs to represent the evolution of modeled entities along time. In Alloy, the common way to do so is to model time by adding a specific set of time instants [Jac06], by giving axioms describing its structure (as traces for instance) and by adding an extra time parameter to every dynamic predicate. This is tedious and cumbersome, if not error-prone.

This shortcoming has long been identified and several propositions [FGPA05, NJ10, VD12] have been made to extend Alloy with facilities for fancier modeling of *behavior*. Still, in all these approaches, the verification remains bounded (because the set of instants is, for instance). [Cun14] makes a step further by implementing a bounded model-checking approach in Alloy allowing time loops. However, up to our knowledge, no Alloy extension leverages a temporal logic, such as LTL for instance, that enjoys a complete decision procedure. The idea of adding temporal logic to FO has been implemented in the tool TLA⁺ [Lam02], where the FO signature is that of ZFC, instead of the arbitrary signatures allowed in Alloy. These remarks led us to study the combination of FO and LTL, in particular to draw questions about the relation between the satisfiability of a FO-LTL formula and the fact that the first-order part of the model is finite. In the literature, the logic FO-LTL has drawn a lot of interest, for decidability questions as well as in database theory [AHdB96]. For instance [HWZ00, HKK⁺03] study decidable fragments, while [Kam68, Mer92] give incompleteness results.

1.2 Contributions

The first question we address here is that of the complexity of satisfiability for FO and FO-LTL when the FO part of the model is bounded (we call this problem BSAT(N) for a given bound N). We are interested in the cost in terms of algorithmic complexity of adding an LTL component to the FO specification language. In Sect. 3, we consider BSAT for FO and FO-LTL depending on whether the quantifier rank (*i.e.* the maximum number of nested quantifiers) of formulas is bounded and whether the bound on the domain is given in unary or binary encoding.

- For pure FO, BSAT(N) is NP-complete if the rank is bounded and if N is given in unary; NEXPTIME-complete otherwise. This case can admittedly be considered *folklore* but it seems it has not been published formally, so we

³ Available at <http://alloy.mit.edu/alloy>.

detail it in the paper for the sake of completeness. We also provide detailed results, showing that formulas of rank 2 with unary predicates are sufficient for NEXPTIME-completeness.

- For FO-LTL, which has been less studied from the point of view of BSAT, we show that this division goes the same except that BSAT is PSPACE-complete in the first case and EXPSpace-complete otherwise (recall that satisfiability for LTL alone is PSPACE-complete [SC85, LP85]). Again, rank 2 formulas with unary predicates are sufficient.

Secondly, since we are only interested in finite models of FO-LTL formulas, it is natural to study which fragments of FO-LTL enjoy the *finite model property* (FMP). Recall that a formula has the FMP if the existence of a model implies the existence of a *finite* one. Many fragments of FO have been shown to enjoy the FMP in the past decades [BGG97, ARS07].

- In Sect. 4.1, we show that any fragment of FO enjoying the FMP (as well as a mild assumption often met in practice) can be “lifted” as a fragment of FO-LTL using also **X** and **F** and still enjoying the FMP (provided the removal of the temporal operators leads back to the original FO fragment).
- We finally show in Sect. 4.2 that with temporal operators **U** or **G**, the FMP is lost, even with strong constraints on the way temporal operators interact with first-order quantifiers.

All these results provide a theoretical insight on the combination of LTL with bounded FO which may be useful in the context of decision procedures based upon SAT or SMT, or in formal methods such as extensions of Alloy or TLA⁺. Another possible application may be in the analysis of *software product lines* [CHS⁺10] where various, related transition systems (which may be described using FO) represent a product family.

2 The Logic FO-LTL

In this section, we define precisely the logic FO-LTL and provide some elements on its expressiveness.

2.1 Syntax

Definition 1 (FO-LTL syntax). *We define the syntax of FO-LTL in the standard way from the following elements (function symbols will also be considered in Sect. 4):*

- a tuple of predicates $\mathcal{P} = (P_1, \dots, P_k)$ (each of which is of any arity) which define relations, between elements of the system, that can vary in time,
- equality = is considered as a particular binary predicate which is static, i.e., its value does not depend on time,
- an infinite countable set Var of variables,

- a finite set $Const$ of constants, representing elements of the system,
- the Boolean connectives $\neg, \vee$,
- the existential quantifier $\exists x$ for each variable $x \in Var$,
- the temporal operators $\mathbf{X}$ (next) and $\mathbf{U}$ (until).

We also add the usual syntactic sugar: $\top, \perp, \wedge, \forall, \Rightarrow, \mathbf{G}, \mathbf{F}, \mathbf{R}$, where

$$F\varphi = \top \mathbf{U} \varphi, \quad G\varphi = \neg(F(\neg\varphi)), \quad \psi \mathbf{R} \varphi = \neg(\neg\varphi \mathbf{U} \neg\psi).$$

Example 1. Let us consider $\mathcal{P} = \{OK, fail\}$, where OK is a nullary predicate, i.e., an atomic proposition, and $fail$ is an unary predicate. We can define the following formula: $\mathbf{G}(\exists x.fail(x) \Rightarrow \mathbf{F} \mathbf{G} \neg OK)$. Intuitively, it expresses that a local bug endangers the whole system and no recovery is possible: if one element of the system fails at some point, then later the system must enter a state where it is not OK and remain in this state forever.

2.2 Semantics

Variables and constants (and more generally terms if we consider functions) are interpreted over a domain D . We consider that the domain and the interpretation of variables and constants do not vary in time. Only the interpretation of predicates can change. The time domain considered throughout the paper is $\mathbb{N}$.

Definition 2 (FO-LTL Structure). An interpretation structure of FO-LTL is a tuple $\mathcal{M} = (D, \sigma_{Const}, \rho)$ where:

- D is the domain,
- $\sigma_{Const} : Const \rightarrow D$ is a valuation for constants,
- $\rho = (P_1^i, \dots, P_k^i)_{i \in \mathbb{N}}$ gives the semantics of each predicate in $\mathcal{P}$ at each instant $i \in \mathbb{N}$. If P_j is a l -ary predicate, then $P_j^i \subseteq D^l$ for each instant $i \in \mathbb{N}$.

We now define the satisfaction of a formula in a structure, in which case the latter is called a *model* of the former.

Definition 3 (Satisfaction Relation). Given a structure $\mathcal{M}$, we inductively define the satisfaction relation $\mathcal{M}, \sigma, i \models \varphi$, where σ maps free variables of φ to elements in D , and $i \in \mathbb{N}$ is the current point in time.

For ease of reading, x and y stand for both variables and constants in this definition. Moreover, we use $\bar{\sigma}$ to denote the interpretation of both variables and constants: $\bar{\sigma}(x) = \sigma(x)$ if $x \in Var$ and $\bar{\sigma}(x) = \sigma_{Const}(x)$ if $x \in Const$.

- $\mathcal{M}, \sigma, i \models x = y$ if $\bar{\sigma}(x) = \bar{\sigma}(y)$
- $\mathcal{M}, \sigma, i \models P_j(x_1, \dots, x_n)$ if $(\bar{\sigma}(x_1), \dots, \bar{\sigma}(x_n)) \in P_j^i$
- $\mathcal{M}, \sigma, i \models \neg\varphi$ if $\mathcal{M}, \sigma, i \not\models \varphi$
- $\mathcal{M}, \sigma, i \models \varphi \vee \psi$ if $\mathcal{M}, \sigma, i \models \varphi$ or $\mathcal{M}, \sigma, i \models \psi$
- $\mathcal{M}, \sigma, i \models \exists x.\varphi$ if there exists $a \in D$ such that $\mathcal{M}, \sigma[x \mapsto a], i \models \varphi$
- $\mathcal{M}, \sigma, i \models \mathbf{X}\varphi$ if $\mathcal{M}, \sigma, i+1 \models \varphi$
- $\mathcal{M}, \sigma, i \models \varphi \mathbf{U} \psi$ if there exists $j \geq i$ such that $\mathcal{M}, \sigma, j \models \psi$, and for all p such that $i \leq p < j$, we have $\mathcal{M}, \sigma, p \models \varphi$

A formula φ without free variables is satisfiable if and only if there exists a structure $\mathcal{M}$ such that $\mathcal{M}, \emptyset, 0 \models \varphi$, and in this case we just note $\mathcal{M} \models \varphi$. (The semantics with function symbols is defined in a similar straightforward way.)

Notice that FO-LTL can be viewed as a fragment of a first-order logic called 2FO, where quantifiers can range either over D or over time. It was shown in [Kam68] that FO-LTL is strictly less expressive than 2FO, as opposed to the classical result that LTL and FO have the same expressive power over discrete time. Detailed definitions and examples regarding 2FO are provided in Appendix A.

3 Complexity of Bounded Satisfiability

We are interested in a problem occurring in practice, where a formula φ of FO or FO-LTL is given together with a bound N , and we want to check the existence of a model of φ with domain of size at most N . This problem is decidable, but its complexity is an interesting question that, as far as we know, has been overlooked (though the FO case can be considered unpublished folklore). We call this problem BSAT and we investigate its complexity for several variants. As explained earlier, this question is of practical interest given the success of formal methods based upon finite model-finding and considering possible temporal extensions of these.

To analyze the complexity of this problem in different settings, we first recall the usual notion of *(quantifier) rank* [Lib04].

Definition 4 (Quantifier Rank). *The (quantifier) rank of a FO-LTL formula is defined by structural recursion as follows:*

- $\text{rk}(x = y) = \text{rk}(P(x_1, \dots, x_k)) = 0$
- $\text{rk}(\neg\varphi) = \text{rk}(\mathbf{X}\varphi) = \text{rk}(\varphi)$
- $\text{rk}(\varphi \vee \psi) = \text{rk}(\varphi \mathbf{U} \psi) = \max(\text{rk}(\varphi), \text{rk}(\psi))$
- $\text{rk}(\exists x, \varphi) = 1 + \text{rk}(\varphi)$.

We are interested in settings where the rank of formulas is bounded, or on the contrary any formula is allowed as input. Restricting rank to a certain bound is a natural assumption in practice, and allows a finer analysis of the parameterized complexity of the BSAT problem. As is standard practice, we write $\text{FO}[k]$ (resp. $\text{FO-LTL}[k]$) for all FO (resp. FO-LTL) formulae of quantifier rank up to k .

This rank is not to be confused with the alternation depth, which increases only with alternations between $\forall$ and $\exists$ quantifiers (or in our case between $\exists$ and $\neg$). We chose here to use quantifier rank to reflect the limited number of variables specified in real-life examples by users, for instance using tools such as Alloy. Notice that bounding the quantifier rank does not trivialize the problem, because we allow arbitrary signatures (again, similarly to the Alloy syntax). We recall that most complexity results on logical formalisms in the literature are relative to fixed signatures.

The following theorem classifies the complexity of BSAT according to three parameters: FO alone versus FO-LTL, N given in unary or binary, and $\text{rk}(\varphi)$

bounded or unbounded. Some of these results regarding FO may be part of folklore, but we reproduce them here for completeness.

Theorem 1. *We consider BSAT for three parameters: logic, encoding, bound on $\text{rk}(\varphi)$ (ranked). The corresponding complexities are given in the following table (N is the bound on the model size, k the bound on $\text{rk}(\varphi)$):*

	N unary	N binary
$\text{FO}[k]$	<i>NP-complete</i>	<i>NEXPTIME-complete</i>
FO	<i>NEXPTIME-complete (even $N = 2$)</i>	<i>NEXPTIME-complete</i>
$\text{FO-LTL}[k]$	<i>PSPACE-complete</i>	<i>EXSPACE-complete</i>
FO-LTL	<i>EXSPACE-complete (even $N = 2$)</i>	<i>EXSPACE-complete</i>

Proofs are given in the remaining of this very Sect. 3.

3.1 First-Order Logic

Lemma 1. *The $\text{BSAT}(N)$ problem for $\text{FO}[k]$ with N in unary is NP-complete.*

Proof.

Membership in NP We show membership in NP by polynomially reducing the problem to SAT. The reduction is informally described here, see Appendix B.1 for the formal construction.

The input formula φ is turned into a quantifier-free formula φ' where quantifiers have been expanded: $\forall x$ (resp. $\exists x$) is replaced by $\bigwedge_{x \in [1, N]}$ (resp. $\bigvee_{x \in [1, N]}$). Constants are turned into integers in the same way, using an initial disjunction on their possible values.

We then turn φ' into a SAT instance φ'' by replacing every occurrence of predicate $R(\vec{a})$, where $\vec{a}$ is an integer vector, by a Boolean variable $x_{R, \vec{a}}$.

This reduction is polynomial because of the unary encoding of N and the bound on $\text{rk}(\varphi)$, and preserves satisfiability.

NP-hardness We now show that BSAT for unary $\text{FO}[k]$ is NP-hard.

We perform a reduction from SAT: given a SAT instance with variables $x_1, \dots, x_n$, we build an instance of BSAT where $x_1, \dots, x_n$ are predicates of arity 0. We can then ask for the existence of a structure of size 0 (or 1), and this will answer the SAT problem. Since we do not need any quantifier to reduce to SAT, we obtain NP-hardness even if the bound on the rank is 0. $\square$

Lemma 2. *The $\text{BSAT}(N)$ problem for $\text{FO}[k]$ with N in binary is NEXPTIME-complete if $k \geq 2$, even restricted to unary predicates. It is NP-complete for $k = 1$.*

Proof. The proof is only sketched here, the detailed version can be found in Appendix B.2. The idea is to reduce directly from a non-deterministic Turing Machine running in exponential time.

Given such a machine M together with an input word u , we want to build a formula φ of FO[2] describing the run of M over u , such that φ has a model of size at most N if and only if M accepts u in at most N steps. Variables in φ will be used to describe positions of the tape of M as well as time instants in the computation of M . For this, we use unary predicates to encode the bits of the cell position $p(x)$ and time instant $t(x)$ described by an element x of the domain. We additionally use predicates $a(x)$ for a in the alphabet of the machine, and $q(x)$ for q in the state space of the machine to specify the content of the cell $p(x)$ at time $t(x)$. To avoid using formulas of rank 3, we also introduce a predicate $a'(x)$ to say that cell $p(x)$ is labelled a at time $t(x) + 1$. We can express that this encoding is sound, and specify the existence of an accepting run of the machine using a formula φ of rank 2. Since N can be specified in binary, and since $|\varphi|$ is polynomial in the size of M , we can show that φ has a model of size N if and only if M has an accepting run of size exponential (2^{n^k}) in its input of size n .

The fact that the problem is in NEXPTIME is proven similarly as in Lemma 1, and is shown for a more general version of the problem in Lemma 4.

On the contrary, if $k = 1$, we show that any satisfiable formula φ of rank 1 has a model of size at most $|\varphi|$, therefore it is in NP to verify the existence of such a model. NP-hardness follows from Lemma 1. $\square$

Lemma 3. *The BSAT(N) problem for unranked FO with N in unary is NEXPTIME-hard, even for $N = 2$.*

Proof. We show that this case is also NEXPTIME-hard.

As before, let M be a non-deterministic Turing machine running in exponential time 2^{n^k} .

This time, we will use predicates of unbounded arity, and encode positions in the tape using binary code. We will actually need only two elements in the structure, named 0 and 1.

To state that a position of binary encoding $\vec{x}$ is labelled by a letter a (resp. a state q), we will use a predicate $a(\vec{x})$ (resp. $q(\vec{x})$) of arity n^k , where each coordinate of $\vec{x}$ is given as a distinct argument.

To mimic the previous proof, we need to be able to compare 2 positions, using a predicate $\vec{x} < \vec{y}$ of arity $2n^k$. Once this order is axiomatized, the reduction can be done as in the previous case.

Therefore, we will only give the relevant new material here, *i.e.* the axioms for $\leq$ of arity $2n^k$ being a total order. These axioms must all be of polynomial length in n , in order to keep the overall reduction polynomial.

We use $\forall \vec{x}$ as a shorthand for $\forall x_1, \forall x_2, \dots, \forall x_{n^k}$. In this way, it suffices to rewrite the axioms of total order using vectors instead of elements. This keeps the size of axioms polynomial, making it grow only by a factor n^k . Note that this does not guarantee that $\leq$ describes the lexicographic order on vectors, in particular the first position could be any vector, but this is not a problem.

Replacing all variables by vectors in the previous proof yields the required reduction. $\square$

The membership in NEXPTIME will be shown in the proof of Lemma 4.

Lemma 4. *The $BSAT(N)$ problem for unranked FO with N in binary is in NEXPTIME.*

Proof. This result implies NEXPTIME-completeness for 3 variants of the First-Order BSAT problem.

Let $\varphi, \vec{e}$ be the input of the problem, where $\vec{e}$ is a binary encoding of N , so $N = O(2^{|\vec{e}|})$. Let $n = |\varphi| + |\vec{e}|$ be the size of the input, and $r = \text{rk}(\varphi)$. The algorithm from the proof of Lemma 1 can be adapted as follows:

- Guess a structure and write it on the tape: a predicate of arity k takes up to N^k cells, so the operation uses time (and space) $O(2^{nk})$.
- Unfold quantifiers of the formula and check predicates. This operation takes time $O(|\varphi|N^r) = O(2^{nr})$.

Overall, the time complexity is in $O(2^{n(k+r)}) = O(2^{n^2})$, since both k and r are bounded by n .

This ends the proof that the most “difficult” FO case of BSAT still has NEXPTIME complexity. Hardness (even for $N = 2$) follows from Lemma 2. $\square$

3.2 An Algorithm for the BSAT Problem for FO-LTL

We now turn to the BSAT problem for FO-LTL, and describe a generic algorithm that we will use for various settings of the problem.

Lemma 5. *The $BSAT(N)$ problem for FO-LTL is in PSPACE if the rank is bounded and N is given in unary, and in EXPSPACE all three other cases.*

The algorithm consists in trying all sizes up to N , and for each of them expand the formula φ into a LTL formula, then use a PSPACE algorithm for LTL satisfiability.

Definition 5 (Expansion of an FO-LTL Formula). *Let us consider a finite domain D , a finite set of constants $Const$, a valuation $\sigma_{Const} : Const \rightarrow D$, a closed FO-LTL formula φ with constants in $Const$ and predicate symbols $P_1, \dots, P_k$, of arities $\alpha_1, \dots, \alpha_k$ respectively. We define the expansion $\exp(\varphi)$ of φ given the domain D as an LTL formula, using alphabet $A = \{A_i(\vec{a}) \mid 1 \leq i \leq k, \vec{a} \in D^{\alpha_i}\}$ by induction on φ . We assume that φ can use elements of D as constants, and σ_{Const} is extended to these new constants in the natural way.*

$$\begin{aligned}
\exp(a = b) &= \top \text{ if } \sigma_{Const}(a) = \sigma_{Const}(b) \text{ and } \perp \text{ otherwise} \\
\exp(P_i(a_1, \dots, a_k)) &= A_i(\sigma_{Const}(a_1), \dots, \sigma_{Const}(a_k)) \\
\exp(\neg\varphi) &= \neg \exp(\varphi) & \exp(\varphi \vee \psi) &= \exp(\varphi) \vee \exp(\psi) \\
\exp(\mathbf{X}\varphi) &= \mathbf{X} \exp(\varphi) & \exp(\varphi \mathbf{U} \psi) &= \exp(\varphi) \mathbf{U} \exp(\psi) \\
\exp(\exists x, \varphi) &= \bigvee_{a \in D} \exp(\varphi[x \leftarrow a])
\end{aligned}$$

It is easy to show by induction that for any φ and D , we have

$$|\exp(\varphi)| = \Theta(|\varphi| \cdot |D|^{\text{rk}(\varphi)}).$$

We can now adapt the algorithm from Lemma 1 to this new setting. In the case where the rank is bounded, we rewrite the formula to bound arity of predicates if the rank is bounded, and guess a structure of size D of size at most N together with σ_{Const} , using space polynomial in $|D|$ (so exponential in the input N is in binary). We then expand φ into $\text{exp}(\varphi)$, of size $O(|\varphi| \cdot N^{\text{rk}(\varphi)})$.

It remains to decide whether the LTL formula $\text{exp}(\varphi)$ is satisfiable, which can be done using space polynomial in $|\text{exp}(\varphi)|$ [SC85]. Therefore this algorithm uses space $O(|\varphi| \cdot N^{\text{rk}(\varphi)})$. It is in PSPACE if the rank is bounded and N is in unary, and EXPSPACE in the other three cases.

3.3 Completeness Results for FO-LTL BSAT

We now show that this algorithm is optimal, by showing that BSAT for FO-LTL is either PSPACE-hard or EXPSPACE-hard depending on the setting.

Lemma 6. *The $\text{BSAT}(N)$ problem for ranked and unranked FO-LTL with N in binary is EXPSPACE-complete, even for $N = 2$. In the ranked case, the bound must be at least 2. The $\text{BSAT}(N)$ problem for FO-LTL $[k]$ with N in unary is PSPACE-complete.*

Proof. The main idea of the proof is to directly encode the run of a Turing machine using exponential space (polynomial space in the ranked case with N in unary), similarly as in the proof of Lemma 2. The main difference is that we now have additional LTL operators, that allow us to encode computation steps without any bound on the number of time instants. Therefore, the first-order domain D will only be used to encode positions of the tape via unary predicates specifying the bits of the position, and that is why we can now encode runs of machines using exponentially more time than space. The detailed reduction can be found in Appendix B.3. $\square$

Finally, the last case is treated in the following lemma.

Lemma 7. *The $\text{BSAT}(N)$ problem for unranked FO-LTL with N in unary is EXPSPACE-complete.*

Proof. We will show that this case is also EXPSPACE-hard, although we can no longer use an element of the structure for each cell of the Turing machine.

We can reuse ideas from Lemma 3, and encode positions using vectors of bits with predicates of unbounded arities. This time, only positions will be encoded this way, as time can be taken care of by LTL. Thus we can start from a machine where only space is exponentially bounded, and time can be doubly exponential.

The construction is then similar to the one from Lemma 3, and yields a reduction showing that this variant of BSAT is also EXPSPACE-complete, even for structures with only 2 elements. $\square$

Other examples of EXPSPACE-complete problems related to deciding small fragments of FO-LTL can be found in [HKK⁺03].

4 Finite Model Property

Since we are only interested in finite models of FO-LTL formulas, it is natural to study which fragments of FO-LTL enjoy the *finite model property* (FMP). We say that a formula has the FMP if the existence of a model implies the existence of a *finite* model (*i.e.*, with finite first-order domain but still infinite time structure). We also say that a fragment *Frag* of some logic has the FMP if all the formulas from *Frag* have the FMP. Many such fragments of FO were exhibited in the past decades.

Function Symbols In this Section we will enrich the syntax of FO-LTL with function symbols. Each function has an arbitrary arity like a predicate, but yields a *term*, which will be interpreted as an element of the domain, as variables and constants. In this case, the parameters of the predicates (including equality) can be arbitrary terms, built by composing variables and constants with functions. For instance, if x and y are variables, a is a constant, f and g are functions, then $f(x, g(x), a) = g(y)$ is a formula.

Example 2. [BGG97, ARS07] The following fragments of *FO*, named following the notation of [BGG97], have the FMP:

- $[\exists^*\forall^*, all] =$ (Ramsey 1930) the class of all sentences with quantifier prefix $\exists^*\forall^*$ over arbitrary relational vocabulary with equality.
- $[\exists^*\forall\exists^*, all] =$ (Ackermann 1928) the class of all sentences with quantifier prefix $\exists^*\forall\exists^*$ over arbitrary relational vocabulary with equality.
- $[\exists^*\forall^2\exists^*, all]$ (Gödel 1932, Kalmár 1933, Schütte 1934) the class of all sentences with quantifier prefix $\exists^*\forall^2\exists^*$ over arbitrary relational vocabulary without equality.
- $[\exists^*, all, all] =$ (Gurevich 1976) the class of all sentences with quantifier prefix $\exists^*$ over arbitrary vocabulary with equality.
- $[\exists^*\forall, all, (1)] =$ (Grädel 1996) the class of all sentences with quantifier prefix $\exists^*\forall$ over vocabulary that contains one unary function and arbitrary predicate symbols with equality.
- $[all, (\omega), (\omega)]$ (Gurevich 1969, Löb 1967) the class of all sentences with arbitrary quantifier prefix over vocabulary that contains an arbitrary number of unary predicates and unary functions without equality
- FO_2 (Mortimer 1975) the class of all sentences of relational vocabulary that contains two variables and equality.

4.1 Lifting FMP from FO to FO-LTL

In this section, we first present general results that allow to lift the finite model property from FO fragments to their temporal extension with operators **X** and **F**. Then, we focus on two particular fragments: the well known Ramsey fragment, for which the extension can be generalized to full LTL, and a fragment that does not fulfill the hypotheses of our general result, but for which the temporal extension with operators **X** and **F** still has the FMP.

Remark 1. In the following, we will only consider formulas in *negation normal form* (NNF), i.e. where negations have been pushed to the leaves. This means negations can only be applied to predicates. Notice that the syntactic sugar mentioned in Sect. 2.1, in particular the operator **R** (dual of **U**) now becomes necessary to retain full expressiveness.

Definition 6. If *Frag* is a fragment of *FO*, and $OP \subseteq \{\mathbf{X}, \mathbf{F}, \mathbf{G}, \mathbf{U}, \mathbf{R}\}$ is a set of temporal operators, we define the fragment *Frag* + *OP* of *FO-LTL* as the formulas with temporal operators from *OP*, where the formula(s) obtained by removing unary temporal operators and replacing binary ones by $\vee$ or $\wedge$ (indifferently), is in *Frag*.

A General Extension Result for Fragments with the FMP

Definition 7 ((Plus-)Replacement of a Formula). If φ, ψ are *FO*-formulas, we say that ψ is a replacement of φ if ψ can be obtained from φ by replacing predicates and functions, i.e., by allowing different occurrences of the same predicate (resp. function) of φ to become distinct predicates (resp. functions) of same arity in ψ , but distinct predicates (resp. functions) in φ are always mapped to distinct predicates (resp. functions) in ψ .

Additionally, we define the notion of plus-replacement where the new predicates and functions can have increased arity.

For instance $\forall x. \exists y. P(x) \vee Q(y)$ is a replacement of $\forall x. \exists y. P(x) \vee P(y)$. Likewise, the formula $\forall x. \exists y. P(x) \vee Q(y, x)$ is a plus-replacement of $\forall x. \exists y. P(x) \vee P(y)$.

Definition 8 (Stability under (Plus-)Replacement). We say that a fragment *Frag* of *FO* with FMP is stable under replacement (resp. plus-replacement) if for all $\varphi \in \text{Frag}$ and for all replacement (resp. plus-replacement) ψ of φ , we have that ψ has the FMP.

In practice, many fragments with FMP considered in the literature are stable under (plus-)replacement. This is for example the case for most of the fragments from Example 2 (see Corollary 1).

Theorem 2 (Frag + **X).** Let *Frag* be a fragment of *FO* with the FMP, stable under replacement. Then the fragment *Frag* + **X** of *FO-LTL* has the FMP.

The proof of this theorem is presented in Appendix C.2. The following theorem, along the same lines, allows more temporal operators but has the stronger assumption of plus-replacement.

Theorem 3 (Frag + **{X, F}).** Let *Frag* be a fragment of *FO* with FMP, stable under plus-replacement. Then *Frag* + **{X, F}** also has the FMP.

Proof. Let φ be a satisfiable formula of *Frag* + **{X, F}**. Let V be the set of variables used in φ and $\{\mathbf{F}_j, j \in J\}$ be the set of **F**-operators in φ , for some finite labeling set $J = \{1, 2, \dots, |J|\}$ such that if $\mathbf{F}_j$ is under the scope of $\mathbf{F}_i$ then $i < j$.

For $\vec{x}$ a list of variables from V , $j \in J \cup \{0\}$, $k \in \mathbb{N}$, and θ a subformula of φ , we define $\llbracket \theta \rrbracket_{\vec{x}}^j$ inductively as follows:

$$\begin{aligned}
\llbracket P(\vec{y}) \rrbracket_{\vec{x}}^{j,k} &= P_{j,k}(\vec{y}, \vec{x}) && \text{variables can appear in both } \vec{y} \text{ and } \vec{x} \\
\llbracket f(\vec{y}) \rrbracket_{\vec{x}}^{j,k} &= f_{j,k}(\vec{y}, \vec{x}) && \text{variables can appear in both } \vec{y} \text{ and } \vec{x} \\
\llbracket \exists z. \theta(\vec{y}) \rrbracket_{\vec{x}}^{j,k} &= \exists z \llbracket \theta(\vec{y}) \rrbracket_{\vec{x}}^{j,k} && \llbracket \forall z. \theta(\vec{y}) \rrbracket_{\vec{x}}^{j,k} = \forall z \llbracket \theta(\vec{y}) \rrbracket_{\vec{x},z}^{j,k} \\
\llbracket \theta(\vec{y}) \wedge \theta'(\vec{y}') \rrbracket_{\vec{x}}^{j,k} &= \llbracket \theta(\vec{y}) \rrbracket_{\vec{x}}^{j,k} \wedge \llbracket \theta'(\vec{y}') \rrbracket_{\vec{x}}^{j,k} \\
\llbracket \theta(\vec{y}) \vee \theta'(\vec{y}') \rrbracket_{\vec{x}}^{j,k} &= \llbracket \theta(\vec{y}) \rrbracket_{\vec{x}}^{j,k} \vee \llbracket \theta'(\vec{y}') \rrbracket_{\vec{x}}^{j,k} \\
\llbracket \mathbf{X} \theta(\vec{y}) \rrbracket_{\vec{x}}^{j,k} &= \llbracket \theta(\vec{y}) \rrbracket_{\vec{x}}^{j,k+1} && \llbracket \mathbf{F}_{j'} \theta(\vec{y}) \rrbracket_{\vec{x}}^{j,k} = \llbracket \theta(\vec{y}) \rrbracket_{\vec{x}}^{j',0}
\end{aligned}$$

To sum up, we index predicates and functions by the label j of the innermost occurrence of $\mathbf{F}$ that has it in its scope, as well as the number k of nested $\mathbf{X}$ since this occurrence. We also add all universally quantified variables as arguments. We additionally remove $\mathbf{F}$ and $\mathbf{X}$ operators in the process.

Let $\psi = \llbracket \varphi \rrbracket_{\emptyset}^{0,0}$. We show that ψ is satisfiable. Let $\mathcal{M} = (D, \rho)$ be a model of φ . This means that for each subformula $\mathbf{F}_j \theta(\vec{y})$ of φ under universally quantified variables $\vec{x}$, there is a function $t_j : D^{|\vec{x}|} \rightarrow \mathbb{N}$ such that $\theta(\vec{y})$ is true at time $t_j(\vec{x})$. We build a model of ψ by setting the value of $P_{j,k}(\vec{y}, \vec{x})$ to $P(\vec{y})$ at time $t_j(\vec{x}) + k$, where $\vec{x}$ is the list of new arguments of P_j (and same with functions). It is straightforward to verify that this is indeed a model of ψ .

Let φ' be φ where the $\mathbf{F}$'s and $\mathbf{X}$'s have been removed, by definition we have $\varphi' \in \text{Frag}$. Since ψ is a plus-replacement of φ' and Frag is stable under plus-replacement, we have ψ has the FMP. Since ψ is satisfiable, there exists a finite model M_f of ψ . Finally, we build from M_f a finite model of φ . For this, we have to choose new values for the $t_j(\vec{x})$, so that no conflicts occur: if $(j, k, \vec{x}) \neq (j', k', \vec{x}')$, then $t_j(\vec{x}) + k \neq t_{j'}(\vec{x}') + k'$. Let K be the maximal number of nested $\mathbf{X}$ (not interleaved with $\mathbf{F}$), and $(\vec{x}_i)_{0 \leq i \leq R}$ be an ordering of all possible vectors $\vec{x}$. We choose $t_j(\vec{x}) = (K+1) \times (R^j + i)$, in order to satisfy the injectivity condition: for all $j, j' \in [0, |J|]$, $k, k' \in [0, K]$, and $\vec{x}, \vec{x}' \in \{\vec{x}_i \mid 0 \leq i \leq R\}$, we have $t_j(\vec{x}) + k = t_{j'}(\vec{x}') + k'$ if and only if $(j, k, \vec{x}) = (j', k', \vec{x}')$. Notice moreover that we respect the condition that if $\mathbf{F}_j \varphi_j$ is a subformula of $\mathbf{F}_i \varphi_i$, then $j > i$ and thus for any value of $\vec{x}, \vec{y}$, we have $t_j(\vec{x}) > t_i(\vec{y})$.

Finally, we build a finite model of φ by setting the value of $P(\vec{y})$ (resp. $f(\vec{y})$) at time i to $P_{j,k}(\vec{y}, \vec{x})$ (resp. $f_{j,k}(\vec{y}, \vec{x})$) if $i = t_j(\vec{x}) + k$ for some $j, k, \vec{x}$, and choosing any values for other time instants.

So φ has a finite model and therefore $\text{Frag} + \{\mathbf{X}, \mathbf{F}\}$ has the FMP. $\square$

Remark 2. It is enough to consider plus-replacement where new arguments are only quantified universally, which is a weaker condition.

Corollary 1. *The following FO-LTL fragments, extending FO fragments mentioned in Example 2, have the FMP:*

$$\begin{aligned}
[\exists^* \forall^*, \text{all}]_+ &= \{\mathbf{X}, \mathbf{F}\} && [\exists^* \forall^* \exists^*, \text{all}]_+ = \{\mathbf{X}, \mathbf{F}\} \\
[\exists^* \forall^2 \exists^*, \text{all}]_+ &= \{\mathbf{X}, \mathbf{F}\} && [\exists^*, \text{all}, \text{all}]_+ = \{\mathbf{X}, \mathbf{F}\} \\
\text{FO}_2 + \{\mathbf{X}, \mathbf{F}\} &&&
\end{aligned}$$

Specific Extensions for Two Fragments In this section, we focus on two fragments of FO: a fragment for which our general result (Theorem 3) does not apply and a fragment for which our general result can be extended to full LTL.

The FMP of the following fragment, even if it is not stable under plus-replacement, can be lifted to its temporal extension with $\mathbf{X}$ and $\mathbf{F}$.

Theorem 4. $[all, (\omega), (\omega)] + \{\mathbf{X}, \mathbf{F}\}$ has the FMP.

Proof. We show that any formula of this fragment has the FMP. We proceed by induction on the number of nested $\mathbf{F}$. The induction hypothesis is actually stronger than the FMP: we show by induction that for such a formula φ , if there is a model then there is a model M with finite domain D and a finite set of time instants T such that M only “looks at T ”, i.e. changing the values of predicates and functions outside of T does not change the truth value of φ .

We start with the base case where there is no $\mathbf{F}$. By Theorem 2 (and its proof), and since $[all, (\omega), (\omega)]$ is stable under replacement (even though it is not stable under plus-replacement), if φ has a model it has a finite one where only the values on a finite set of instants matter.

We now turn to the induction step, and consider a formula φ with $n + 1$ nested $\mathbf{F}$. By considering the outermost occurrences of $\mathbf{F}$, the formula φ can be written $\varphi'(\mathbf{F}\varphi_1, \mathbf{F}\varphi_2, \dots, \mathbf{F}\varphi_k)$, where φ' contains no $\mathbf{F}$ but may contain quantifiers, and for every $i \in [1, k]$, φ_i has at most n nested $\mathbf{F}$ and may contain free variables.

We assume that φ has a model M , and without loss of generality we note j the index in $[1, k]$ such that $\mathbf{F}\varphi_i$ is true in M (for at least one valuation of its free variables) if and only if $i \leq j$. This means in particular that for all $i \leq j$, φ_i has a model. By the induction hypothesis, for all $i \leq j$, φ_i can be set to true in a model M_i with a finite domain D_i and that only looks at a finite set of instants T_i .

Moreover, $\varphi'' = \varphi'(\top, \dots, \top, \perp, \dots, \perp)$ (with j times $\top$) is satisfiable, and by the base case has a model M' with a finite domain D' that only looks at a finite set of instants T' (that will be used as the first instants of the model).

We now build a model M_f for φ with a finite domain D , that we define as a set of cardinality $\max(|D'|, |D_1|, |D_2|, \dots, |D_j|)$.

We define a sequence of time instants $(t_i)_{1 \leq i \leq j}$ such that at time t_i the formula φ_i is true for a particular valuation of its free variables, and at $t_i + |T_i|$, it is true for another valuation of its free variables that are universally quantified, and so on, until we have considered all the possible values in D for these universally quantified variables. So we define the t_i inductively as follows : $t_1 = |T'|$ and $t_i = t_{i-1} + |T_{i-1}| \times |D|^r$, where r is the number of nested universal quantifiers in φ' .

We now describe the predicates and function values in M_f . At times $t \in [0, t_1 - 1]$, we mimic the model M' . This gives the value of predicates and functions for $|D'|$ elements of D . All the remaining elements can be set to behave as any element of D' , for instance the first one. Since equality cannot be tested, and predicates and functions are monadic, the truth value of φ'' is preserved.

For all $i \in [1, j]$, we use M_i to fix the valuation of predicates and functions at times $[t_i, t_i + |T_i| - 1]$. Then, from $t_i + |T_i|$, we consider another possible assignment of universally quantified variables and define the valuation of predicates and functions accordingly. This way, we obtain a model of φ_i starting at time t_i , and therefore a model of $\mathbf{F}\varphi_i$ starting at time 0.

Since φ' is monotonous in its arguments (no $\mathbf{F}$ can be under a negation), and we preserved the value $\top$ for all $\mathbf{F}\varphi_i$ with $i \leq j$, the truth value of φ on M_f is that of φ' , which is true thanks to the valuations on $[0, t_0]$.

We have therefore built a model M_f of φ with finite domain, and only looking at a finite number of time instants. $\square$

The result of Ramsey that the $\exists^*\forall^*$ fragment has the FMP is generalized in the following theorem. See Appendix C.1 for a proof, omitted here due to space constraints.

Theorem 5. *We consider here FO-LTL without function symbols. Let $\varphi = \exists x_1 \dots \exists x_n. \psi(x_1, \dots, x_n)$, where ψ is a FO-LTL formula without any $\exists$ quantifiers. Then if φ is satisfiable, it has a model with domain of size at most $n + c$, where c is the number of constants in the vocabulary.*

4.2 Axioms of Infinity using LTL

We now give examples showing that adding LTL to fragments of FO with the FMP allows to write axioms of infinity, therefore losing the FMP. This holds even when strong restrictions are enforced on the way LTL operators interact with first-order quantifiers.

Extending the Ramsey Fragment. First, let us remark that the constraint from Theorem 5 that existential quantifiers are not under the scope of temporal operators is necessary, as showed by the following formula which is only satisfiable by infinite models, using a unary predicate P :

$$\mathbf{G}(\exists x. P(x) \wedge \mathbf{XG} \neg P(x)).$$

Indeed, it is straightforward to show that a different x_n is needed to satisfy the formula at each different time instant $n \in \mathbb{N}$, as the condition on the predicate P guarantees that the same x can never be used twice.

Separating Quantifiers and LTL. We now give examples where a fragment of FO loses the FMP when extended with LTL, even without nesting quantifiers under temporal operators.

The following FO-LTL formula is an axiom of infinity with a $\forall\exists$ quantifier prefix, and where no first-order quantifiers are under the scope of LTL operators. It uses one constant c and one unary predicate P :

$$\forall x \exists y. P(c) \wedge \mathbf{G}(P(x) \Rightarrow \mathbf{X}(P(y) \wedge \mathbf{G} \neg P(x))).$$

This sentence only has infinite models, as the predicate P must be true on a different element at each instant of time. However, as recalled in Example 2, in FO without LTL, if only one quantifier $\forall$ is used the FMP is guaranteed (or alternatively, formulas with two variables also have the FMP).

This example can actually be replaced using $\mathbf{U}$ instead of $\mathbf{G}$, showing that it suffices to be able to refer to an “unbounded” (as opposed to “infinite”) number of time instants to force models to be infinite, as showed by the following example:

$$\forall x \exists y. P(x) \wedge ((P(x) \wedge P(y)) \mathbf{U} (\neg P(x) \wedge P(y))).$$

This time, we used values of the predicate P in the past instead of the future to guarantee that the same x cannot be used twice.

5 Conclusion

Motivated by the possible extension of formal methods based upon finite model finding (such as Alloy or various decision procedures based upon SAT or SMT techniques) with temporal reasoning, we have investigated FO-LTL with finite FO domains in two ways: (1) we studied the complexity of the satisfiability for FO-LTL (and FO alone) when the FO part of the model is bounded; (2) we studied cases where we can lift the FMP of FO fragments to their temporal extensions.

Several question are still open. On the complexity side, it remains to settle the case of FO-LTL[1] with N in binary. Related to the FMP, even if we showed in Sect. 4.1 that for a particular FO fragment that is not stable under plus-replacement, the FMP can still be lifted to its temporal extension with operators $\mathbf{X}$ and $\mathbf{F}$, it is not clear whether this assumption can be dropped in Theorem 3. Another open question is whether we can find a reasonable condition under which we can extend an FO fragment with temporal operators $\mathbf{G}$ or $\mathbf{U}$ without losing the FMP. Indeed, these operators bring an expressiveness that is very useful in practice but we showed in Sect. 4.2 that they behave badly with respect to the FMP.

References

- [AHdB96] Serge Abiteboul, Laurent Herr, and Jan Van den Bussche. Temporal versus first-order logic to query temporal databases. In *Proceedings of the Fifteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, June 3-5, 1996, Montreal, Canada*, pages 49–57, 1996.
- [ARS07] Aharon Abadi, Alexander Moshe Rabinovich, and Mooly Sagiv. Decidable fragments of many-sorted logic. In *14th International Conference on Logic for Programming, Artificial Intelligence, and Reasoning (LPAR 2007)*, pages 17–31, 2007.
- [BGG97] Egon Börger, Erich Grädel, and Yuri Gurevich. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer, 1997.

- [BKMJ15] Hamid Bagheri, Eunsuk Kang, Sam Malek, and Daniel Jackson. Detection of design flaws in the Android permission protocol through bounded verification. In *FM 2015*, 2015.
- [CHS⁺10] Andreas Classen, Patrick Heymans, Pierre-Yves Schobbens, Axel Legay, and Jean-François Raskin. Model checking lots of systems: efficient verification of temporal properties in software product lines. In *ICSE 2010*, pages 335–344. ACM, 2010.
- [Cun14] Alcino Cunha. Bounded model checking of temporal formulas with alloy. In *ABZ 2014*, volume 8477 of *Lecture Notes in Computer Science*, pages 303–308. Springer Berlin Heidelberg, 2014.
- [FGPA05] Marcelo F. Frias, Juan P. Galeotti, Carlos López Pombo, and Nazareno Aguirre. DynAlloy: upgrading Alloy with actions. In *ICSE 2005*, pages 442–451, 2005.
- [HKK⁺03] Ian M. Hodkinson, Roman Kontchakov, Agi Kurucz, Frank Wolter, and Michael Zakharyashev. On the computational complexity of decidable fragments of first-order linear temporal logics. In *TIME-ICTL 2003*, pages 91–98, 2003.
- [Hod99] Ian M. Hodkinson. Note on games in temporal logics, 1999. Lecture notes for LUATCS meeting, Johannesburg.
- [HWZ00] Ian M. Hodkinson, Frank Wolter, and Michael Zakharyashev. Decidable fragments of first-order temporal logics. *Annals of Pure and Applied Logic*, 106(1–3):85 – 134, 2000.
- [Jac06] Daniel Jackson. *Software Abstractions - Logic, Language, and Analysis*. MIT Press, 2006.
- [Kam68] Hans W. Kamp. *Tense Logic and the Theory of Linear Order*. Phd thesis, University of Warsaw, 1968.
- [Lam02] Leslie Lamport. *Specifying Systems, The TLA+ Language and Tools for Hardware and Software Engineers*. Addison-Wesley, 2002.
- [Lib04] Leonid Libkin. *Elements of Finite Model Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer, 2004.
- [LP85] Orna Lichtenstein and Amir Pnueli. Checking that finite state concurrent programs satisfy their linear specification. In *Proceedings of the 12th ACM SIGACT-SIGPLAN symposium on Principles of programming languages*, pages 97–107. ACM, 1985.
- [Mer92] Stephan Merz. Decidability and incompleteness results for first-order temporal logics of linear time. *Journal of Applied Non-Classical Logics*, 2(2):139–156, 1992.
- [NJ10] Joseph P. Near and Daniel Jackson. An imperative extension to Alloy. In *ABZ 2010*, pages 118–131, 2010.
- [NRZ⁺15] Chris Newcombe, Tim Rath, Fan Zhang, Bogdan Munteanu, Marc Brooker, and Michael Deardeuff. How Amazon Web Services uses formal methods. *Commun. ACM*, 58(4):66–73, 2015.
- [SC85] A. Prasad Sistla and Edmund M. Clarke. The complexity of propositional linear temporal logics. *J. ACM*, 32(3):733–749, 1985.
- [TJ07] Emina Torlak and Daniel Jackson. Kodkod: A relational model finder. In *TACAS 2007*, volume 4424 of *LNCS*, pages 632–647. Springer, 2007.
- [VD12] Amirhossein Vakili and Nancy A. Day. Temporal logic model checking in Alloy. In *ABZ 2012*, pages 150–163, 2012.
- [Zav12] Pamela Zave. Using lightweight modeling to understand Chord. *SIG-COMM Comput. Commun. Rev.*, 42(2):49–57, 2012.

Appendix

A Expressiveness of FO-LTL

FO-LTL allows to express properties of evolving systems. But this could also be done using FO with two types of variables: one type ranging over the domain (the domain variables) and one ranging over time (the temporal variables).

Definition 9 (Syntax of 2FO). *The syntax of 2FO (for 2-sorted FO) is defined in the standard way from the following elements:*

- a finite set of predicates. Each predicate takes as arguments an arbitrary number of domain variables and one temporal variable specifying at which instant we want to evaluate the predicate.
- the equality predicate $=$, which is a binary predicate on domain variables
- the ordering predicate $\leq$, which is a binary predicate on temporal variables
- the boolean connectives $\neg, \vee$,
- the existential quantifier $\exists$ over domain variables
- the existential quantifier $\exists^T$ over temporal variables

Lemma 8. [Kam68] *Every FO-LTL formula can be translated into an equivalent 2FO formula.*

Proof. Straightforward, by induction on the FO-LTL formula, expressing temporal operators with first-order formula. $\square$

Kamp showed [Kam68] that 2FO and FO-LTL do not have same expressive power, even if there is only one unary predicate P . Namely the 2FO formula

$$\exists^T t, t'. t < t' \wedge \forall x. P(x, t) \Leftrightarrow P(x, t')$$

cannot be expressed in FO-LTL. This formula expresses that the truth value of predicate P is exactly identical at two different instants of time.

Remark 3. This counter-example requires an infinite first-order domain (*i.e.* relative to x), as it trivializes to true on finite domains with infinite time structure. We can adapt it to make it work even if the semantics is restricted to finite first-order domains (using $+1$ as syntactic sugar):

$$\exists^T t, t'. t < t' \wedge \forall x. P(x, t) \Leftrightarrow P(x, t') \Leftrightarrow P(x, t' + 1).$$

However, it is showed in [Hod99] that FO-LTL is complete for the “simple” fragment of 2FO where every subformula of the form $\exists x. \varphi$ has at most one free time variable.

B Complexity of BSAT

B.1 NP Membership for Lemma 1

We want to show that given a FO-formula φ of arity at most R , and an integer $N \in \mathbb{N}$ given in unary, it is in NP to decide whether there exists a model of φ of size at most N .

We will first convert the formula φ to a formula φ' without quantifiers, and where predicates of arity k range over $[1, N]^k$.

We first remove constants by replacing them with existentially quantified variables, and use the formula $\exists c_1, \exists c_2, \dots, \exists c_p. \varphi$ where $c_1, \dots, c_p$ are the constants used in φ .

We compute a SAT instance φ' from this formula, with Boolean variables of the form $x_{i,\vec{a}}$ where P_i is a relation symbol of arity k used in φ , and $\vec{a}$ is a vector of k integers from $[1, N]$.

This is done inductively via the operator sat defined as follows:

- $sat(a = b) = \top$ if $a = b$ and $\perp$ otherwise, where a, b are integers from $[1, N]$.
- $sat(P_i(\vec{a}) = x_{i,\vec{a}})$, where $\vec{a}$ is a vector of integers from $[1, N]$.
- $sat(\neg\phi) = \neg sat(\phi)$,
- $sat(\phi \vee \psi) = sat(\phi) \vee sat(\psi)$,
- $sat(\phi \wedge \psi) = sat(\phi) \wedge sat(\psi)$,
- $sat(\exists x, \phi) = \bigvee_{a \in [1, N]} sat(\phi[x \leftarrow a])$.
- $sat(\forall x, \phi) = \bigwedge_{a \in [1, N]} sat(\phi[x \leftarrow a])$.

It is straightforward to show that satisfiability by a model of size N is preserved by this translation, by interpreting the value of $x_{i,\vec{a}}$ as $P_i(\vec{a})$.

Moreover, the bound on $rk(\varphi)$ as well as the unary encoding of N guarantees that the size of $sat(\varphi)$ is polynomial in $|\varphi|$. Therefore, it suffices compute the SAT-instance $sat(\varphi)$ in polynomial time, and decide in NP whether it is satisfiable. This yields a NP algorithm for this variant of BSAT.

B.2 Proof of Lemma 2

NEXPTIME-Hardness Proof of Lemma 2 We show NEXPTIME-hardness by reducing directly from a non-deterministic Turing machine M running in exponential time 2^{n^k} , for a certain k , where n is the size of the input of M . Notice that a tape of size 2^{n^k} is enough for all computations. In order to simplify a bit the construction, we assume that the tape of M is infinite in only one direction, and is labeled by $\mathbb{N}$. Given M and a word u of size n , we want to determine whether $u \in L(M)$, by reducing it polynomially to an instance of this variant of BSAT.

Let A be the alphabet of the Turing Machine M including a special blank letter $B \in A$, Q its set of states, and Δ its transition table. Let q_{init}, q_{fin} be the initial and final states of M , respectively.

Elements of the domain D will serve to encode both time instant and positions on the tape. This is done via $2n^k$ unary predicates $T_1, \dots, T_{n^k}, P_1, \dots, P_{n^k}$:

for an element $x \in D$, using the convention that false stands for 0 and true for 1, the binary number $T_1(x)T_2(x) \dots T_{n^k}(x)$ is the time instant $t(x)$ represented by x and $P_1(x)P_2(x) \dots P_{n^k}(x)$ is the cell position $p(x)$ represented by x .

We will encode configurations of M via predicates $a(x)$ for all $a \in A$ (respectively $q(x)$ for all $q \in Q$), meaning that the cell $p(x)$ at time $t(x)$ is labelled with the letter a (respectively contains the head of the machine in state q). In order to achieve the reduction with only formulas of rank 2, we also add for each $a \in A$ a predicate $a'(x)$, meaning that $p(x)$ is labelled by a at instant $t(x+1)$.

We will use the following quantifier-free auxiliary formulas to manipulate binary encodings:

- $\varphi_{t=}(x, y) = \bigwedge_{1 \leq i \leq n^k} T_i(x) \Leftrightarrow T_i(y)$, meaning $t(x) = t(y)$.
- $\varphi_{t+1}(x, y) = \bigvee_{1 \leq i \leq n^k} \left(\neg T_i(x) \wedge T_i(y) \right) \wedge \bigwedge_{i < j \leq n^k} (T_j(x) \wedge \neg T_j(y)) \wedge \bigwedge_{1 \leq j < i} (T_j(x) \Leftrightarrow T_j(y))$, meaning $t(x) + 1 = t(y)$.
- $\varphi_{p=}(x, y)$ and $\varphi_{p+1}(x, y)$ are defined similarly, by replacing T_i 's with P_i 's.
- $first_t(x) = \bigwedge_{1 \leq i \leq n^k} \neg T_i(x)$.
- $first_p(x)$, $last_t(x)$, $last_p(x)$ are defined in the same way.
- $Q(x) = \bigvee_{q \in Q} q(x)$

We include the following axioms describing the encoding, gathered as a conjunction in a formula φ_{ax} :

- First element: $\exists x. first_t(x) \wedge first_p(x)$.
- Each element (but the last) has a successor on one component while preserving the other:
 $\forall x. \bigwedge_{\{\alpha, \beta\}=\{p, t\}} (last_\alpha(x) \vee (\exists y. \varphi_{\alpha+1}(x, y) \wedge \varphi_{\beta=}(x, y)))$.
- At most one state at any time: $\forall x. \forall y. (Q(x) \wedge Q(y) \wedge \varphi_{t=}(x, y)) \Rightarrow x = y$.
- At least one state at any time: $\forall x. \exists y. \varphi_{t=}(x, y) \wedge Q(y)$.
- At most one state on each cell: $\forall x. \bigwedge_{p \neq q \in Q} \neg p(x) \vee \neg q(x)$.
- One letter on each cell:
 $\forall x. \bigvee_{a \in A} a(x) \wedge \bigwedge_{a \neq b \in A} \neg a(x) \vee \neg b(x)$.
- Consistency of a' : $\forall x. \forall y. \varphi_{t+1}(x, y) \Rightarrow \bigwedge_{a \in A} a'(x) \Leftrightarrow a(y)$.

Notice that all these formulas have rank at most 2, therefore $\text{rk}(\varphi_{ax}) = 2$

The initial configuration must start with $u = a_0 a_1 \dots a_{n-1}$, with the machine in the initial state on the first position. We write $\varphi_{p=i}(x)$ the formula making explicit that $p(x) = i$ and $t(x) = 0$ by listing all values of the predicates T_i and P_i , and $\varphi_{p \geq n}$ is defined similarly, for $p(x) \geq n$ and $t(x) = 0$. These formulas are quantifier-free.

The initial configuration is specified by the rank 1 formula.

$$\varphi_{in} = (\exists x. \varphi_{p=0}(x) \wedge q_{init}(x)) \wedge (\forall x. (\bigwedge_{0 \leq i \leq n-1} \varphi_{p=i}(x) \Rightarrow a_i(x)) \wedge \forall x. (\varphi_{p \geq n}(x) \Rightarrow B(x))).$$

The fact that the machine eventually halts and accepts is described with $\varphi_{fin} = \exists x.q_{fin}(x)$.

Finally, let $\Delta \subseteq Q \times A \times \{\leftarrow, \rightarrow\} \times A \times Q$ be the transition table of M , we use the following rank 2 formula to describe the temporal evolution according to Δ :

$$\begin{aligned} \varphi_{trans} := & \forall x.(last_t(x)) \vee \\ & \bigwedge_{q_1 \in Q} \left\{ q_1(x) \Rightarrow \exists y. \left[\varphi_{t+1}(x, y) \right. \right. \\ & \quad \wedge \bigvee_{(q_1, a, \leftarrow, b, q_2) \in \Delta} (a(x) \wedge b'(x) \wedge q_2(y) \wedge \varphi_{p+1}(y, x)) \\ & \quad \vee \bigvee_{(q_1, a, \rightarrow, b, q_2) \in \Delta} (a(x) \wedge b'(x) \wedge q_2(y) \wedge \varphi_{p+1}(x, y)) \\ & \quad \left. \left. \wedge (\forall y. (\varphi_{t=}(x, y) \wedge \neg \varphi_{p=}(x, y)) \Rightarrow \bigvee_{a \in A} a(y) \wedge a'(y)) \right] \right\} \end{aligned}$$

Finally, let $\psi = \varphi_{ax} \wedge \varphi_{in} \wedge \varphi_{fin} \wedge \varphi_{trans}$.

It is straightforward to verify that ψ has a model of size (at most) $N = 2^{n^k}$ if and only if u is accepted by M . In this binary version of BSAT, N can be specified using only n^k bits, so polynomially with respect to $|M| + |u|$. Since $|\psi|$ is polynomial in u , uses only unary predicates, and $\text{rk}(\psi) = 2$, this shows that the ranked FO-fragment of BSAT with N in binary is NEXPTIME-hard even if limited to unary predicates (provided the bound for the rank is at least 2).

NP-completeness Proof of Lemma 2 for rank 1 We now turn to the case where we are limited to rank 1 formulas, and show that it is NP-complete.

NP-hardness is directly obtained by reduction from SAT: a SAT instance can be rewritten as a rank 0 formula with one constant and n unary predicates (or no constant and a rank 1 formula).

We now show NP membership. We first remark that it is useless to have predicates with arity more than one, since they can only be called with several occurrences of the same variable, and possibly constants, so it is equivalent to a list of unary predicates. For instance if we have a binary predicate R and a constant a , we need three binary predicates: $R_1(x) = R(x, x)$, $R_2(x) = R(x, a)$, and $R_3(x) = R(a, x)$ and axioms guaranteeing consistency: $R_1(a) = R_2(a) = R_3(a)$.

A formula ψ of rank 1 is therefore equivalent to a positive Boolean combination of formulas of the form $\forall x.\varphi(x)$ and $\exists x.\varphi(x)$ where the φ are Boolean combinations of unary predicates. Assume there is a model M of ψ , we show there is a model M' of size at most $|\psi|$. Indeed, consider the existential formulas φ_i from the Boolean combination that are true in M , for each of them there is an element x_i witnessing it. It suffices to take M' to be the restriction of M to the union of these x_i . Due to the structure of the formula, each component of the

positive Boolean combination that was true in M is still true in M' , therefore M' is a model of ψ .

This shows that it is enough to verify the existence of a linear-size model for ψ , thus the problem is in NP.

B.3 Proof of Lemma 6

We will show EXPSPACE-hardness for the cases where N is given in binary. As in the FO case, we reduce directly from a Turing machine M running in exponential space 2^{n^k} , with unidirectional tape. Given M and a word u of size n , we want to determine whether $u \in L(M)$, by reducing it polynomially to an instance of BSAT for FO-LTL with N binary.

Let A be the alphabet of the Turing Machine M including a special blank letter $B \in A$, Q its set of states, and Δ its transition table. Let q_i and q_f be the initial and final states of M , respectively.

The FO structure we will work with is the tape of the machine, of size $N = 2^{n^k}$ (it takes only n^k bits to specify N), with a unary predicate for each element of A and Q , as well as a list of unary predicates $P_1, \dots, P_{n^k}$ specifying the bits of the position on the tape. A difference with the proof of Lemma 2 is that now an element x represents only a position $p(x)$. We will use the implicit time of LTL to describe the evolution of the computation. As a consequence we do not need predicates $a'(x)$ of the previous proof, as they can be expressed with $\mathbf{X}a(x)$ instead.

We use auxiliary formulas $\varphi_{p=}$, $\varphi_{p+1}(x, y)$, $first_p(x)$, $last_p(x)$ and $Q(X)$ identical to the one in the proof of Lemma 2.

We define a formula φ_{ax} as in the proof of Lemma 2, to enforce the semantics of the predicates. We now have to manage time with LTL instead of a binary encoding, which gives the following list of conjuncts for φ_{ax} :

- Predicates P_i do not change over time: $\forall x. \mathbf{G}(\bigwedge_{1 \leq i \leq n^k} (P_i(x) \Leftrightarrow \mathbf{X}P_i(x)))$.
- First element: $\exists x. first_p(x)$.
- Each element (but the last) has a successor:
 $\forall x. (last_p(x) \vee (\exists y. \varphi_{p+1}(x, y)))$.
- At most one state at any time: $\mathbf{G}(\forall x. \forall y. (Q(x) \wedge Q(y)) \Rightarrow x = y)$.
- At least one state at any time: $\mathbf{G}(\exists x. Q(x))$.
- At most one state on each cell: $\mathbf{G}(\forall x. \bigwedge_{p \neq q \in Q} \neg p(x) \vee \neg q(x))$.
- One letter on each cell:
 $\mathbf{G}(\forall x. \bigvee_{a \in A} a(x) \wedge \bigwedge_{a \neq b \in A} \neg a(x) \vee \neg b(x))$.

The initial configuration is specified as in the proof of Lemma 2, omitting the specification $t = 0$.

The final configuration is described with $\varphi_{fin} = \exists x. q_{fin}(x)$.

LTL formulas can be used to describe the evolution of the tape: we can say in FO-LTL that at most two positions can change between two consecutive instants, and they have to do it according to the transition table $\Delta \subseteq Q \times A \times \{\leftarrow, \rightarrow\} \times A \times Q$ of M .

This is done in the following formula:

$$\begin{aligned} \varphi_{trans} := \exists x. \Big[& \left(\bigvee_{(p,a,\leftarrow,b,q) \in \Delta} \exists y. (\varphi_{p+1}(y,x) \wedge p(x) \wedge a(x) \wedge \mathbf{X}(b(x) \wedge q(y))) \right) \\ & \vee \bigvee_{(p,a,\rightarrow,b,q) \in \Delta} \exists y. (\varphi_{p+1}(x,y) \wedge p(x) \wedge a(x) \wedge \mathbf{X}(b(x) \wedge q(y))) \Big) \\ & \wedge \forall y. ((y = x) \vee \bigvee_{a \in A} a(y) \wedge \mathbf{X}(a(y))) \Big]. \end{aligned}$$

Finally, let $\psi = \varphi_{ax} \wedge \varphi_{in} \wedge (\varphi_{trans} \mathbf{U} \varphi_{fin})$. It is straightforward to verify that ψ has a model of size (at most) N if and only if u is accepted by M . Since $|\psi|$ is polynomial in $|\varphi| + \log_2(N)$, and has rank 2, this shows both the ranked (with bound at least 2) and unranked variants of BSAT for FO-LTL with binary N given are EXPSpace-hard, and therefore EXPSpace-complete by Lemma 5.

The same proof gives PSPACE-hardness if N is in unary, and therefore PSPACE-completeness for the ranked unary version. Indeed, the exact same construction works, except now N must remain polynomial in the size of the input, and therefore we must start with a machine running in polynomial space in order to be able to reduce it to BSAT.

C Finite Model Property

C.1 Proof of Theorem 5

Let φ be a subformula of FO-LTL using constants $Const$, free variables V and without $\exists$ quantifiers. Let $\mathcal{M} = (D, \sigma, \rho)$ be a model of φ where $\sigma : Const \cup V \rightarrow D$. We also assume that negations have been pushed to the leaves of φ , and that there is no functional symbol. This means that we need the dual connectives $\wedge$ of $\vee$ and $\mathbf{R}$ of $\mathbf{U}$ in the syntax.

We show by induction on φ that restricting the domain D to $D' = \sigma(Const \cup V)$ and ρ accordingly by restricting predicates to D' yields another model of φ .

- If φ is atomic (equality or predicate), then only free variables and constants, referring to elements from D' are used.
- If $\varphi = \varphi_1 \vee \varphi_2$ or $\varphi = \varphi_1 \wedge \varphi_2$, we conclude by induction on φ_1 and φ_2 .
- If $\varphi = \forall x. \psi(x)$, we use the induction hypothesis over ψ : for any $a \in D$ we can restrict the domain to $\sigma(Const \cup V) \cup \{a\}$ and make $\psi(a)$ true. This means that if we restrict the domain to $D' = \sigma(Const \cup V)$, it makes φ true, since the $\forall$ quantifier will only constrain the $a \in D'$.
- $\varphi = \mathbf{X} \varphi_1$, $\varphi = \varphi_1 \mathbf{U} \varphi_2$, or $\varphi = \varphi_1 \mathbf{U} \varphi_2$, we can conclude by induction on φ_1, φ_2 and by the definition of the semantics of FO-LTL that restricting the domain to D' at all times yields a valid model.

C.2 Proof of Theorem 2

Let φ be a satisfiable formula of $\text{Frag} + \mathbf{X}$, and n be the number of nested $\mathbf{X}$ operators in φ . We build ψ by replacing in φ every occurrence of predicate P by P_i , where $i \in [0, n]$ is the number of $\mathbf{X}$ operators having this occurrence of P in their scope; and by removing all $\mathbf{X}$ operators. Notice that ψ is satisfiable, because any model of φ can be turned into a model of ψ , by setting the value of P_i and f_i to the value of P, f at time i .

Let φ' be φ where the $\mathbf{X}$'s have been removed: by definition we have $\varphi' \in \text{Frag}$. Since ψ is a replacement of φ' and Frag is stable under replacement, we have that ψ enjoys the FMP. Since ψ is satisfiable, there exists a finite model M of ψ . Finally, we can turn M into a finite model of φ , by setting the value of P (resp. f) at time i to P_i (resp. f_i), and choosing any values for time instants strictly larger than n , that do not affect the formula. We conclude that φ has a finite model, and therefore $\text{Frag} + \mathbf{X}$ has the FMP. $\square$