“Most of” leads to undecidability: Failure of adding frequencies to LTL

Bartosz Bednarczyk^{1,2} and Jakub Michaliszyn²

¹ Computational Logic Group, Technische Universität Dresden, Dresden, Germany

² Institute of Computer Science, University of Wrocław, Wrocław, Poland
{bartosz.bednarczyk, jakub.michaliszyn}@cs.uni.wroc.pl

Abstract. Linear Temporal Logic (LTL) interpreted on finite traces is a robust specification framework popular in formal verification. However, despite the high interest in the logic in recent years, the topic of their quantitative extensions is not yet fully explored. The main goal of this work is to study the effect of adding weak forms of percentage constraints (e.g. that *most of* the positions in the past satisfy a given condition, or that σ is the *most-frequent* letter occurring in the past) to fragments of LTL. Such extensions could potentially be used for the verification of influence networks or statistical reasoning. Unfortunately, as we prove in the paper, it turns out that percentage extensions of even tiny fragments of LTL have undecidable satisfiability and model-checking problems. Our undecidability proofs not only sharpen most of the undecidability results on logics with arithmetics interpreted on words known from the literature, but also are fairly simple. We also show that the undecidability can be avoided by restricting the allowed usage of the negation, and discuss how the undecidability results transfer to first-order logic on words.

1 Introduction

Linear Temporal Logic [29] (LTL) interpreted on finite traces is a robust logical framework used in formal verification [1,18,19]. However, LTL is not perfect: it can express whether some event happens or not, but it cannot provide any insight on how frequently such an event occurs or for how long such an event took place. In many practical applications, such *quantitative* information is important: think of optimising a server based on how frequently it receives messages or optimising energy consumption knowing for how long a system is usually used in rush hours. Nevertheless, there is a solution: one can achieve such goals by adding quantitative features to LTL.

It is known that adding quantitative operators to LTL often leads to undecidability. The proofs, however, typically involve operators such as “next” or “until”, and are often quite complicated (see the discussion on the related work below). In this work, we study the logic $\text{LTL}_{\mathbf{F}}$, a fragment of LTL where the only allowed temporal operator is “sometimes in the future” $\mathbf{F}$. We extend its language with two types of operators, sharing a similar “percentage” flavour: with the *Past-Majority* $\mathbf{PM}\varphi$ operator (stating that most of the past positions satisfy

a formula φ), and with the *Most-Frequent-Letter* **MFL** σ predicates (meaning that the letter σ is among the most frequent letters appearing in the past). These operators can be used to express a number of interesting properties, such as *if a process failed to enter the critical section, then the other process was in the critical section the majority of time*. Of course, for practical applications, we could also consider richer languages, such as parametrised versions of these operators, *e.g.* stating that *at least a fraction p of positions in the past satisfies a formula*. However, we show, as our main result, that even these very simple percentage operators raise undecidability when combined with **F**.

To make the undecidability proof for both operators similar, we define an intermediate operator, **Half**, which is satisfied when exactly half of the past positions satisfy a given formula. The **Half** operator can be expressed easily with **PM**, but not with **MFL** — we show, however, that we can simulate it to an extent enough to show the undecidability. Our proof method relies on enforcing a model to be in the language $(\{wht\}\{shdw\})^+$, for some letters *wht* and *shdw*, which a priori seems to be impossible without the “next” operator. Then, thanks to the specific shape of the models, we show that one can “transfer” the truth of certain formulae from positions into their successors, hence the “next” operator can be partially expressed. With a combination of these two ideas, we show that it is possible to write equicardinality statements in the logic. Finally, we perform a reduction from the reachability problem of Two-counter Machines [26]. In the reduction, the equicardinality statements will be responsible for handling zero-tests. The idea of transferring predicates from each position into its successor will be used for switching the machine into its next configuration.

The presented undecidability proof of LTL with percentage operators can be adjusted to extensions of fragments of first-order logic on finite words. We show that $\text{FO}_M^2[<]$, *i.e.* the two-variable fragment of first-order logic admitting the majority quantifier **M** and linear order predicate $<$ has an undecidable satisfiability problem. Here the meaning of a formula $\text{M}x.\varphi(x, y)$ is that at least a half of possible interpretations of x satisfies $\varphi(x, y)$. Our result sharpens an existing undecidability proof for (full) FO with Majority from [23] (since in our case the number of variables is limited) but also $\text{FO}^2[<, succ]$ with arithmetics from [25] (since our counting mechanism is weaker and the successor relation *succ* is disallowed). On the positive side, we show that the undecidability heavily depends on the presence of the negation in front of the percentage operators. To do so, we introduce a logic, extending the full LTL, in which the usage of percentage operators is possible, but suitably restricted. For this logic, we show that the satisfiability problem is decidable.

All the above-mentioned results can be easily extended to the model checking problem, where the question is whether a given Kripke structure satisfies a given formula. The full version of the paper is available on arXiv [4].

1.1 Related work

The first paper studying the addition of quantitative features to logic was [21], where the authors proved undecidability of Weak MSO with Cardinalities. They

also developed a model of so-called Parikh Automaton, a finite automaton imposing a semi-linear constraint on the set of its final configurations. Such an automaton was successfully used to decide logics with counting as well as logics on data words [27,17]. Its expressiveness was studied in [11].

Another idea in the realm of quantitative features is availability languages [20], which extend regular expressions by numerical occurrence constraints on the letters. However, their high expressivity leads to undecidable emptiness problems. Weak forms of arithmetics have also attracted interest from researchers working on temporal logics. Several extensions of LTL were studied, including extensions with counting [24], periodicity constraints [14], accumulative values [7], discounting [2], averaging [9] and frequency constraints [8]. A lot of work was done to understand LTL with timed constraints, *e.g.* a metric LTL was considered in [28]. However, its complexity is high and its extensions are undecidable [3].

Arithmetical constraints can also be added to the First-Order logic (FO) on words via so-called counting quantifiers. It is known that weak MSO on words is decidable with threshold counting and modulo-counting (thanks to the famous Büchi theorem [10]), while even FO on words with percentage quantifiers becomes undecidable [23]. Extensions of fragments of FO on words are often decidable, *e.g.* the two-variable fragment FO^2 with counting [12] or FO^2 with modulo-counting [25]. The investigation of decidable extensions of FO^2 is limited by the undecidability of FO^2 on words with Presburger constraints [25].

Among the above-mentioned logics, the formalisms of this paper are most similar to Frequency LTL [8]. The satisfiability problem for Frequency LTL was claimed to be undecidable, but the undecidability proof as presented in [8] is bugged (see [9, Sec. 8] for discussion). It was mentioned in [9] that the undecidability proof from [8] can be patched, but no correction was published so far. Our paper not only provides a valid proof but also sharpens the result, as we use a way less expressive language (*e.g.* we are allowed to use neither the “until” operator nor the “next” operator). We also believe that our proof is simpler. The second-closest formalism to ours is average-LTL [9]. The main difference is that the averages of average-LTL are computed based on the future, while in our paper, the averages are based on the past. The second difference, as in the previous case, is that their undecidability proof uses more expressive operators, such as the “until” operator.

2 Preliminaries

We recall definitions concerning logics on words and temporal logics (*cf.* [15]).

Words and logics. Let AP be a countably-infinite set of *atomic propositions*, called here also *letters*. A finite *word* $\mathbf{w} \in (2^{\text{AP}})^*$ is a non-empty finite sequence of *positions* labelled with sets of letters from AP . A set of words is called a *language*. Given a word $\mathbf{w}$, we denote its i -th position with $\mathbf{w}_i$ (where the first position is $\mathbf{w}_0$) and its prefix up to the i -th position with $\mathbf{w}_{\leq i}$. We usually use the letters p, q, i, j to denote positions. With $|\mathbf{w}|$ we denote the length of $\mathbf{w}$.

The syntax of $\text{LTL}_{\mathbf{F}}$, a fragment of LTL with only the *finally* operator $\mathbf{F}$, is defined with the grammar: $\varphi, \varphi' ::= a \text{ (with } a \in \text{AP}) \mid \neg\varphi \mid \varphi \wedge \varphi' \mid \mathbf{F}\varphi$.

The satisfaction relation $\models$ is defined for words as follows:

$$\begin{aligned} \mathfrak{w}, i &\models a && \text{if } a \in \mathfrak{w}_i \\ \mathfrak{w}, i &\models \neg\varphi && \text{if not } \mathfrak{w}, i \models \varphi \\ \mathfrak{w}, i &\models \varphi_1 \wedge \varphi_2 && \text{if } \mathfrak{w}, i \models \varphi_1 \text{ and } \mathfrak{w}, i \models \varphi_2 \\ \mathfrak{w}, i &\models \mathbf{F}\varphi && \text{if } \exists j \text{ such that } |\mathfrak{w}| > j \geq i \text{ and } \mathfrak{w}, j \models \varphi. \end{aligned}$$

We write $\mathfrak{w} \models \varphi$ if $\mathfrak{w}, 0 \models \varphi$. The usual Boolean connectives: $\top, \perp, \vee, \rightarrow, \leftrightarrow$ can be defined, hence we will use them as abbreviations. Additionally, we use the *globally* operator $\mathbf{G}\varphi := \neg\mathbf{F}\neg\varphi$ to speak about events happening globally in the future.

Percentage extension. In our investigation, *percentage operators* **PM**, **MFL** and **Half** are added to $\text{LTL}_{\mathbf{F}}$.

The operator **PM** φ (read as: *majority in the past*) is satisfied if at least half of the positions in the past satisfy φ :

$$\mathfrak{w}, i \models \mathbf{PM}\varphi \text{ if } |\{j < i : \mathfrak{w}, j \models \varphi\}| \geq \frac{i}{2}$$

For example, the formula $\mathbf{G}(r \leftrightarrow \neg g) \wedge \mathbf{G}\mathbf{PM}r \wedge \mathbf{G}\mathbf{F}(g \wedge \mathbf{PM}g)$ is true over words where each *request* r is eventually fulfilled by a *grant* g , and where each grant corresponds to at least one request. This can be also seen as the language of balanced parentheses, showing that with the operator **PM** one can define properties that are not regular.

The operator **MFL** σ (read as: *most-frequent letter in the past*), for $\sigma \in \text{AP}$, is satisfied if σ is among the letters with the highest number of appearances in the past, *i.e.*

$$\mathfrak{w}, i \models \mathbf{MFL}\sigma \text{ if } \forall \tau \in \text{AP}. |\{j < i : \mathfrak{w}, j \models \sigma\}| \geq |\{j < i : \mathfrak{w}, j \models \tau\}|$$

For example, the formula $\mathbf{G}\neg(r \wedge g) \wedge \mathbf{G}\mathbf{MFL}r \wedge \mathbf{G}\mathbf{F}(g \wedge \mathbf{MFL}g)$ again defines words where each request is eventually fulfilled, but this time the formula allows for states where nothing happens (*i.e.* when both r and g are false).

The last operator, **Half** is used to simplify the forthcoming undecidability proofs. This operator can be satisfied only at even positions, and its intended meaning is *exactly half of the past positions satisfy a given formula*.

$$\mathfrak{w}, i \models \mathbf{Half}\varphi \text{ if } |\{j < i : \mathfrak{w}, j \models \varphi\}| = \frac{i}{2}$$

It is not difficult to see that the operator **Half** φ can be defined in terms of the past-majority operator as $\mathbf{PM}(\varphi) \wedge \mathbf{PM}(\neg\varphi)$ and that **Half** φ can be satisfied only at even positions.

In the next sections, we distinguish different logics by enumerating the allowed operators in the subscripts, *e.g.* $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$ or $\text{LTL}_{\mathbf{F}, \mathbf{MFL}}$.

Computational problems Kripke structures are commonly used in verification to formalise abstract models. A Kripke structure is composed of a finite set S of *states*, a set of *initial* states $I \subseteq S$, a total *transition* relation $R \subseteq S \times S$, and a finite *labelling function* $\ell : S \rightarrow 2^{\text{AP}}$. A *trace* of a Kripke structure is a finite word

$\ell(s_0), \ell(s_1), \dots, \ell(s_k)$ for any $s_0, s_1, \dots, s_k$ satisfying $s_0 \in I$ and $(s_i, s_{i+1}) \in R$ for all $i < k$.

The *model-checking problem* amounts to checking whether *some* trace of a given Kripke structure satisfies a given formula φ . In the *satisfiability problem*, or simply in *SAT*, we check whether an input formula φ has a *model*, i.e. a finite word $\mathbf{w}$ witnessing $\mathbf{w} \models \varphi$.

3 Playing with Half Operator

Before we jump into the encoding of Minsky machines, we present some exercises to help the reader understand the expressive power of the logic $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$. The tools established in the exercises play a vital role in the undecidability proofs provided in the following section.

We start from the definition of shadowy words.

Definition 1. Let *wht* and *shdw* be fixed distinct atomic propositions from AP . A word $\mathbf{w}$ is *shadowy* if its length is even, all even positions of $\mathbf{w}$ are labelled with *wht*, all odd positions of $\mathbf{w}$ are labelled with *shdw*, and no position is labelled with both letters.

We will call the positions satisfying *wht* simply *white* and their successors satisfying *shdw* simply their *shadows*.

The following exercise is simple in LTL, but becomes much more challenging without the $\mathbf{X}$ operator.

Exercise 1. There is an $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ formula ψ_{shadowy} defining shadowy words.

Solution. We start with the “base” formula $\varphi_{\text{init}}^{\text{exl}} := \text{wht} \wedge \mathbf{G}(\text{wht} \leftrightarrow \neg \text{shdw}) \wedge \mathbf{G}(\text{wht} \rightarrow \mathbf{F} \text{shdw})$, which states that the position 0 is labelled with *wht*, each position is labelled with exactly one letter among *wht*, *shdw* and that every white eventually sees a shadow in the future. What remains to be done is to ensure that only odd positions are shadows and that only even positions are white.

In order to do that, we employ the formula $\varphi_{\text{odd}}^{\text{exl}} := \mathbf{G}((\mathbf{Half} \text{ wht}) \leftrightarrow \text{wht})$. Since $\mathbf{Half}$ is never satisfied at odd positions, the formula $\varphi_{\text{odd}}^{\text{exl}}$ stipulates that odd positions are labelled with *shdw*. An inductive argument shows that all the even positions are labelled with *wht*: for the position 0, it follows from $\varphi_{\text{init}}^{\text{exl}}$. For an even position $p > 0$, assuming (inductively) that all even positions are labelled with *wht*, the formula $\varphi_{\text{odd}}^{\text{exl}}$ ensures that p is labelled with *wht*.

Putting it all together, the formula $\psi_{\text{shadowy}} := \varphi_{\text{init}}^{\text{exl}} \wedge \varphi_{\text{odd}}^{\text{exl}}$ is as required. $\square$

In the next exercise, we show that it is possible to transfer the presence of certain letters from white positions into their shadows. It justifies the usage of “shadows” in the paper.

We introduce the so-called *counting terms*. For a formula φ , word $\mathbf{w}$ and a position p , by $\#_{\varphi}^<(\mathbf{w}, p)$ we denote the total number of positions among $0, \dots, p-1$

satisfying φ , *i.e.* the size of $\{p' < p \mid \mathbf{w}, p' \models \varphi\}$. We omit $\mathbf{w}$ in counting terms if it is known from the context.

Exercise 2. Let σ and $\tilde{\sigma}$ be distinct letters from $\text{AP} \setminus \{wht, shdw\}$. There is an $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ formula $\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}}$, such that $\mathbf{w} \models \varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}}$ iff:

1. $\mathbf{w}$ is shadowy,
2. only white (resp., shadow) positions of $\mathbf{w}$ can be labelled σ (resp., $\tilde{\sigma}$) and
3. for any even position p we have: $\mathbf{w}, p \models \sigma \Leftrightarrow \mathbf{w}, p+1 \models \tilde{\sigma}$.

Solution. Note that the first two conditions can be expressed with the conjunction of ψ_{shadowy} , $\mathbf{G}(\sigma \rightarrow wht)$ and $\mathbf{G}(\tilde{\sigma} \rightarrow shdw)$. The last condition is more involving. Assuming that the words under consideration satisfy conditions 1–2, it is easy to see that the third condition is equivalent to expressing that all white positions p satisfy the equation ($\heartsuit$):

$$(\heartsuit) : \#_{wht \wedge \sigma}^<(\mathbf{w}, p) = \#_{shdw \wedge \tilde{\sigma}}^<(\mathbf{w}, p)$$

supplemented with the condition ($\diamond$), ensuring that the last white position satisfies the condition 3, *i.e.*

$$(\diamond) : \text{ for the last white position } p \text{ we have: } \mathbf{w}, p \models \sigma \Leftrightarrow \mathbf{w}, p+1 \models \tilde{\sigma}.$$

The proof of the following lemma can be found in the appendix.

Lemma 1. *Let $\mathbf{w}$ be a word satisfying the conditions 1–2. Then $\mathbf{w}$ satisfies the condition 3 iff $\mathbf{w}$ satisfies ($\diamond$) and for all white positions p the equation ($\heartsuit$) holds.*

Going back to Exercise 2, we show how to define ($\heartsuit$) and ($\diamond$) in $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$, taking advantage of shadowness of the intended models. Take an arbitrary white position p of $\mathbf{w}$. The equation ($\heartsuit$) for p is clearly equivalent to:

$$(\heartsuit') : \#_{wht \wedge \sigma}^<(\mathbf{w}, p) + \left(\frac{p}{2} - \#_{shdw \wedge \tilde{\sigma}}^<(\mathbf{w}, p) \right) = \frac{p}{2}$$

Since p is even, we infer that $\frac{p}{2} \in \mathbb{N}$. From the shadowness of $\mathbf{w}$, we know that there are exactly $\frac{p}{2}$ shadows in the past of p . Moreover, each shadow satisfies either $\tilde{\sigma}$ or $\neg\tilde{\sigma}$. Hence, the expression $\frac{p}{2} - \#_{shdw \wedge \tilde{\sigma}}^<(\mathbf{w}, p)$ from ($\heartsuit'$), can be replaced with $\#_{shdw \wedge \neg\tilde{\sigma}}^<(\mathbf{w}, p)$. Finally, since *wht* and *shdw* label disjoint positions, the property that every white position p satisfies ($\heartsuit$) can be written as an $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ formula $\varphi_{(\heartsuit)} := \mathbf{G}(wht \rightarrow \mathbf{Half}([wht \wedge \sigma] \vee [shdw \wedge \neg\tilde{\sigma}]))$. Its correctness follows from the correctness of each arithmetic transformation and the semantics of $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$.

For the property ($\diamond$), we first need to define formulae detecting the last and the second to last positions of the model. Detecting the last position is easy: since the last position of $\mathbf{w}$ is shadow, it is sufficient to express that it sees only

shadows in its future, i.e. $\varphi_{last}^{ex2} := \mathbf{G}(shdw)$. Similarly, a position is second to last if it is white and it sees only white or last positions in the future, which results in a formula $\varphi_{stl}^{ex2} := wht \wedge \mathbf{G}(wht \vee \varphi_{last}^{ex2})$. Note that the correctness of φ_{last}^{ex2} and φ_{stl}^{ex2} follows immediately from shadowness. Hence, we can define the formula $\varphi_{(\diamond)}$ as $\mathbf{F}(\varphi_{stl}^{ex2} \wedge \sigma) \leftrightarrow \mathbf{F}(\varphi_{last}^{ex2} \wedge \tilde{\sigma})$. The conjunction of $\varphi_{(\heartsuit)}$ and $\varphi_{(\diamond)}$ formulae gives us to $\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{trans}$. $\square$

We consider a generalisation of shadowy models, where each shadow mimics all letters from a finite set $\Sigma \subseteq \text{AP}$ rather than just a single letter σ . Such a generalisation is described below. In what follows, we always assume that for each $\sigma \in \Sigma$ there is a unique $\tilde{\sigma}$, which is different from σ , and $\tilde{\sigma} \notin \Sigma$. Moreover, we always assume that $\sigma_1 \neq \sigma_2$ implies $\tilde{\sigma}_1 \neq \tilde{\sigma}_2$.

Definition 2. Let $\Sigma \subseteq \text{AP} \setminus \{wht, shdw\}$ be a finite set. A shadowy word $\mathbf{w}$ is called truly Σ -shadowy, if for every letter $\sigma \in \Sigma$ only the white (resp. shadow) positions of $\mathbf{w}$ can be labelled with σ (resp. $\tilde{\sigma}$) and every white position p of $\mathbf{w}$ satisfies $\mathbf{w}, p \models \sigma \Leftrightarrow \mathbf{w}, p+1 \models \tilde{\sigma}$.

Knowing the solution for the previous exercise, it is easy to come up with a formula $\psi_{shadowy}^{truly-\Sigma}$ defining truly Σ -shadowy models: just take the conjunction of $\psi_{shadowy}$ and $\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{trans}$ over all letters $\sigma \in \Sigma$. The correctness follows immediately from Exercise 2.

Corollary 1. The formula $\psi_{shadowy}^{truly-\Sigma}$ defines the language of truly Σ -shadowy words.

The next exercise shows how to compare cardinalities in $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ over truly Σ -shadowy models. We are not going to introduce any novel techniques here, but the exercise is of great importance: it is used in the next section to encode zero tests of Minsky machines.

Exercise 3. Let Σ be a finite subset of $\text{AP} \setminus \{wht, shdw\}$ and let $\alpha \neq \beta \in \Sigma$. There exists an $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ formula $\psi_{\# \alpha = \# \beta}$ such that for any truly Σ -shadowy word $\mathbf{w}$ and any of its white positions p : the equivalence $\mathbf{w}, p \models \psi_{\# \alpha = \# \beta} \Leftrightarrow \#_{wht \wedge \alpha}^<(\mathbf{w}, p) = \#_{wht \wedge \beta}^<(\mathbf{w}, p)$ holds.

The solution is in the appendix, here we briefly discuss the main idea. Follow the previous exercise. The main difficulty is to express the equality of counting terms, written as $\text{LHS} = \text{RHS}$. Note that it is clearly equivalent to $\text{LHS} + (\frac{p}{2} - \text{RHS}) = \frac{p}{2}$. Unfold $\frac{p}{2}$ on the left hand side, *i.e.* replace it with the total number of shadows in the past. Use the fact that $\mathbf{w}$ satisfies $\varphi_{\sigma \rightsquigarrow \bar{\sigma}}^{\text{trans}}$, which implies the equality $\#_{\text{wht} \wedge \beta}^<(\mathbf{w}, p) = \#_{\text{shdw} \wedge \bar{\beta}}^<(\mathbf{w}, p)$. Finally, get rid of subtraction and write an $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ formula by employing **Half**. The presented exercises show that the expressive power of $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ is so high that, under a mild assumption of truly-shadowness, it allows us to perform cardinality comparison. We are now only a step away from showing undecidability of the logic, which is tackled next.

4 Undecidability of LTL extensions

This section is dedicated to the main technical contribution of the paper, namely that $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$, $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$ and $\text{LTL}_{\mathbf{F}, \mathbf{MFL}}$ have undecidable satisfiability and model checking problems. We start from $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$. Then, the undecidability of $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$ will follow immediately from the fact that **Half** is definable by **PM**. Finally, we will show how the undecidability proof can be adjusted to $\text{LTL}_{\mathbf{F}, \mathbf{MFL}}$.

We start by recalling the basics on Minsky Machines.

Minsky machines A *deterministic Minsky machine* is, roughly speaking, a finite transition system equipped with two unbounded-size natural counters, where each counter can be incremented, decremented (only in the case it is positive), and tested for being zero. Formally, a Minsky machine $\mathcal{A}$ is composed of a finite set of *states* Q with a distinguished *initial* state q_0 and a transition function $\delta : (Q \times \{0, +\}^2) \rightarrow (\{-1, 0, 1\}^2 \times (Q \setminus \{q_0\}))$ satisfying three additional requirements: whenever $\delta(q, f, s) = (\bar{f}, \bar{s}, q')$ holds, $\bar{f} = -1$ implies $f = +$, $\bar{s} = -1$ implies $s = +$ (*i.e.* it means that only the positive counters can be decremented) and $q \neq q'$ (the machine cannot enter the same state two times in a row). Intuitively, the first coordinate of δ describes the current state of the machine, the second and the third coordinates tell us whether the current value of the i -th counter is zero or positive, the next two coordinates denote the update on the counters and the last coordinate denotes the target state.

We define a *run* of a Minsky machine $\mathcal{A}$ as a sequence of consecutive transitions of $\mathcal{A}$. Formally, a run of $\mathcal{A}$ is a finite word $\mathbf{w} \in (Q \times \{0, +\}^2 \times \{-1, 0, 1\}^2 \times Q \setminus \{q_0\})^+$ such that, when denoting $\mathbf{w}_i$ as $(q^i, f^i, s^i, \bar{f}^i, \bar{s}^i, q_N^i)$, all the following conditions are satisfied:

1. $q^0 = q_0$ and $f^0 = s^0 = 0$,
2. for each i we have $\delta(q^i, f^i, s^i) = (\bar{f}^i, \bar{s}^i, q_N^i)$,
3. for each $i < |\mathbf{w}|$ we have $q_N^i = q^{i+1}$,
4. for each i , f^i equals 0 iff $f^0 + \dots + f^{i-1} = 0$, and $+$ otherwise; similarly s^i is 0 if $\bar{s}^0 + \dots + \bar{s}^{i-1} = 0$ and $+$ otherwise.

It is not hard to see that this definition is equivalent to the classical one [26]. We say that a Minsky machine *reaches* a state $q \in Q$ if there is a run with a letter containing q on its last coordinate. It is well known that the problem of checking whether a given Minsky machine reaches a given state is undecidable [26].

4.1 “Half of” meets the halting problem

We start from presenting the overview of the claimed reduction. Until the end of Section 4, let us fix a Minsky machine $\mathcal{A} = (Q, q_0, \delta)$ and its state $\mathbf{q} \in Q$. Our ultimate goal is to define an $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ formula $\psi_{\mathcal{A}}^{\mathbf{q}}$ such that $\psi_{\mathcal{A}}^{\mathbf{q}}$ has a model iff $\mathcal{A}$ reaches $\mathbf{q}$. To do so, we define a formula $\psi_{\mathcal{A}}$ such that there is a one-to-one correspondence between the models of $\psi_{\mathcal{A}}$ and runs of $\mathcal{A}$. Expressing the reachability of q , and thus $\psi_{\mathcal{A}}^{\mathbf{q}}$, based on $\psi_{\mathcal{A}}$ is easy.

Intuitively, the formula $\psi_{\mathcal{A}}$ describes a shadowy word $\mathfrak{w}$ encoding on its white positions the consecutive letters of a run of $\mathcal{A}$. In order to express it, we introduce a set $\Sigma_{\mathcal{A}}$, composed of the following distinguished atomic propositions:

- $from_q$ and to_q for all states $q \in Q$,
- $fVal_c$ and $sVal_c$ for counter values $c \in \{0, +\}$, and
- fOP_{op} and sOP_{op} for all operations $op \in \{-1, 0, 1\}$.

We formalise the one-to-one correspondence as the function run , which takes an appropriately defined shadowy model and returns a corresponding run of $\mathcal{A}$. More precisely, the function $run(\mathfrak{w})$ returns a run whose i th configuration is $(q, f, s, \bar{f}, \bar{s}, q_N)$ if and only if the i th white configuration of $\mathfrak{w}$ is labelled with $from_q, fVal_f, sVal_s, fOP_{\bar{f}}, sOP_{\bar{s}}$ and to_{q_N} .

The formula $\psi_{\mathcal{A}}$ ensures that its models are truly $\Sigma_{\mathcal{A}}$ -shadowy words representing a run satisfying properties P1–P4. To construct it, we start from $\psi_{shadowy}^{truly-\Sigma_{\mathcal{A}}}$ and extending it with four conjuncts. The first two of them represent properties P1–P2 of runs. They can be written in $\text{LTL}_{\mathbf{F}}$ in an obvious way.

To ensure the satisfaction of the property P3, we observe that in some sense the letters $from_q$ and to_q are paired in a model, *i.e.* always after reaching a state in $\mathcal{A}$ you need to get out of it (the initial state is an exception here, but we assumed that there are no transitions to the initial state). Thus, to identify for which q we should set the $from_q$ letter on the position p , it is sufficient to see for which state we do not have a corresponding pair, *i.e.* for which state q the number of white $from_q$ to the left of p is not equal to the number of white to_q to the left of p . We achieve this in the spirit of Exercise 3.

Finally, the satisfaction of the property P4 can be achieved by checking for each position p whether the number of white fOP_{+1} to the left of p is the same as the number of white fOP_{-1} to the left of p , and similarly for the second counter. This reduces to checking an equicardinality of certain sets, which can be done by employing shadows and Exercise 3.

The reduction Now we are ready to present the claimed reduction.

We first restrict the class of models under consideration to truly $\Sigma_{\mathcal{A}}$ -shadowy words (for the feasibility of equicardinality encoding) with a formula $\psi_{shadowy}^{truly-\Sigma_{\mathcal{A}}}$. Then, we express that the models satisfy properties P1 and P2. The first property can be expressed with $\psi_{P1} := from_{q_0} \wedge fVal_0 \wedge sVal_0$.

The property P2 will be a conjunction of two formulae. The first one, namely ψ_{P2}^1 , is an immediate implementation of P2. The second one, *i.e.* ψ_{P2}^2 , is not necessary, but simplifies the proof; we require that no position is labelled by more

than six letters from $\Sigma_{\mathcal{A}}$.

$$\begin{aligned}\psi_{P2}^1 &:= \mathbf{G} \left(\text{wht} \rightarrow \bigvee_{\delta(q,f,s)=(\bar{f},\bar{s},q_N)} \text{from}_q \wedge fVal_f \wedge sVal_s \wedge fOP_{\bar{f}} \wedge sOP_{\bar{s}} \wedge to_{q_N} \right), \\ \psi_{P2}^2 &:= \mathbf{G} \bigwedge_{\substack{p_1, \dots, p_7 \in \Sigma_{\mathcal{A}} \\ p_1, \dots, p_7 \text{ are pairwise different}}} \neg(p_1 \wedge p_2 \wedge \dots \wedge p_7).\end{aligned}$$

We put $\psi_{P2} := \psi_{P2}^1 \wedge \psi_{P2}^2$ and $\psi_{enc-basics} := \psi_{shadowy}^{truly-\Sigma_{\mathcal{A}}} \wedge \psi_{P1} \wedge \psi_{P2}$.

We now formalise the correspondence between intended models and runs. Let *run* be the function which takes a word $\mathbf{w}$ satisfying $\psi_{enc-basics}$ and returns the word $\mathbf{w}^A$ such that $|\mathbf{w}^A| = |\mathbf{w}|/2$ and for each position i we have:

$$\begin{aligned}(\Leftarrow) : \mathbf{w}_i^A = (q, f, s, \bar{f}, \bar{s}, q_N) \text{ iff} \\ \mathbf{w}_{2i} \supseteq \{ \text{wht}, \text{from}_q, fVal_f, sVal_s, fOP_{\bar{f}}, sOP_{\bar{s}}, to_{q_N} \}.\end{aligned}$$

The definition of $\psi_{enc-basics}$ makes the function *run* correctly defined and unambiguous, and that the results of *run* satisfy properties P1 and P2.

Fact 5 *The function run is uniquely defined and returns words satisfying P1 and P2.*

What remains to be done is to ensure properties P3 and P4. Both formulas rely on the tools established in Exercise 3 and are defined as follows:

$$\begin{aligned}\psi_{P3} &:= \mathbf{G} \left(\text{wht} \rightarrow \bigwedge_{q \in Q \setminus \{q_0\}} (\text{from}_q \vee \psi_{\# \text{from}_q = \# to_q}) \right). \\ \psi_{P4} &:= \mathbf{G} (fVal_0 \rightarrow \psi_{\# fOP_{+1} = \# fOP_{-1}}) \\ &\quad \wedge \mathbf{G} (sVal_0 \rightarrow \psi_{\# sOP_{+1} = \# sOP_{-1}}) \\ &\quad \wedge \mathbf{G} (\text{wht} \rightarrow (fVal_0 \leftrightarrow \neg fVal_+)) \wedge \mathbf{G} (\text{wht} \rightarrow (sVal_0 \leftrightarrow \neg sVal_+))\end{aligned}$$

Lemma 2. *If $\mathbf{w}$ satisfies $\psi_{enc-basics} \wedge \psi_{P3}$, then $\text{run}(\mathbf{w})$ satisfies P1–P3.*

Proof. The satisfaction of the properties P1 and P2 by $\text{run}(\mathbf{w})$ follows from Fact 5. Ad absurdum, assume that $\text{run}(\mathbf{w})$ does not satisfy P3. It implies the existence of a white position p in $\mathbf{w}$ such that $\mathbf{w}, p \models to_q$ but $\mathbf{w}, p+2 \models \text{from}_{q'}$ for some $q \neq q'$. By our definition of Minsky machines, we conclude that $\mathbf{w}, p \models \text{from}_{q''}$ for some $q'' \neq q$. Thus, $\mathbf{w}, p \not\models \text{from}_q$.

From the satisfaction of ψ_{P3} by $\mathbf{w}$ we know that $\mathbf{w}, p \models \psi_{\# \text{from}_q = \# to_q}$. Let k be the total number of positions labelled with from_q before p . Since $\mathbf{w}, p \models \psi_{\# \text{from}_q = \# to_q}$ holds, by Exercise 3 we infer that the number of positions satisfying to_q before p is also equal to k . Since $\mathbf{w}, p+2 \not\models \text{from}_q$ and from the satisfaction of ψ_{P3} by $\mathbf{w}$ we once more conclude $\mathbf{w}, p+2 \models \psi_{\# \text{from}_q = \# to_q}$. But such a situation clearly cannot happen due to the fact that the number of to_q in the past is equal to $k+1$, while the number of from_q in the past is k . $\square$

Finally, let us define $\psi_{\mathcal{A}}$ as $\psi_{enc-basics} \wedge \psi_{P3} \wedge \psi_{P4}$. The use of $\leftrightarrow$ in ψ_{P4} guarantees that $fVal_0$ labels exactly the white positions having the counter empty (and similarly for the second counter). The counters are never decreased from 0, thus the white positions not satisfying $fVal_0$ are exactly those having the first counter positive.

The proof of the forthcoming fact relies on the correctness of Exercise 3 and is quite similar to the proof of Lemma 2, and is presented in the appendix.

Lemma 3. *If $\mathfrak{w}$ satisfies $\psi_{\mathcal{A}}$, then $run(\mathfrak{w})$ is a run of $\mathcal{A}$.*

Lastly, to show that the encoding is correct, we need to show that each run has a corresponding model. It is again easy: it can be shown by constructing an appropriate $\mathfrak{w}$; the white positions are defined according to $(\rightsquigarrow)$, and the shadows can be constructed accordingly.

Fact 6 *If $\mathfrak{w}^{\mathcal{A}}$ is a run of $\mathcal{A}$, then there is a word $\mathfrak{w} \models \psi_{\mathcal{A}}$ s.t. $run(\mathfrak{w}) = \mathfrak{w}^{\mathcal{A}}$.*

Let $\psi_{\mathcal{A}}^q := \psi_{\mathcal{A}} \wedge \mathbf{F}(to_q)$. Observe that the formula $\psi_{\mathcal{A}}^q$ is satisfiable if and only if $\mathcal{A}$ reaches q . The “if” part follows from Lemma 3 and the satisfaction of the conjunct $\mathbf{F}(to_q)$ from $\psi_{\mathcal{A}}$. The “only if” part follows from Fact 6. Hence, from undecidability of the reachability problem Minsky machines we infer our main theorem:

Theorem 1. *The satisfiability problem for $LTL_{\mathbf{F}, \mathbf{Half}}$ is undecidable.*

6.1 Undecidability of model-checking

For a given alphabet Σ , we can define a Kripke structure $\mathcal{K}_{\Sigma}$ whose set of traces is the language $(2^{\Sigma})^+$: the set of states S of $\mathcal{K}_{\Sigma}$ is composed of all subsets of Σ , all states are initial (*i.e.* $I = S$), the transition relation is the maximal relation ($R = S \times S$) and $\ell(X) = X$ for any subset $X \subseteq \Sigma$. It follows that a formula φ over an alphabet Σ is satisfiable if and only if there is a trace of $\mathcal{K}_{\Sigma}$ satisfying φ . From the undecidability of the satisfiability problem for $LTL_{\mathbf{F}, \mathbf{Half}}$ we get:

Theorem 2. *Model-checking of $LTL_{\mathbf{F}, \mathbf{Half}}$ formulae over Kripke structures is undecidable.*

The decidability can be regained if additional constraints on the shape of Kripke structures are imposed: model-checking of $LTL_{\mathbf{F}, \mathbf{Half}}$ formulae over *flat* structures is decidable [13].

As discussed earlier, the **Half** operator can be expressed in terms of the **PM** operator. Hence, we conclude:

Corollary 2. *Model-checking and satisfiability problems for $LTL_{\mathbf{F}, \mathbf{PM}}$ are undecidable.*

6.2 Most-Frequent Letter and Undecidability

We next turn our attention to the **MFL** operator, which turns out to be a little bit problematic. Typically, formulae depend only on the atomic propositions that they explicitly mentioned. Here, it is not the case. Consider a formula $\varphi = \mathbf{MFL} a$ and words $\mathbf{w}_1 = \{a\}\{\}\{a\}$ and $\mathbf{w}_2 = \{a, b\}\{b\}\{a, b\}$. Clearly, $\mathbf{w}_1, 2 \models \varphi$ whereas $\mathbf{w}_2, 2 \not\models \varphi$. This can be fixed in many ways – for example, by parametrising **MFL** with a domain, so that it expresses that “ a is the most frequent letter among $b_1, \dots, b_n$ ”. We show, however, that even this very basic version of **MFL** is undecidable. The proof is an adaptation of our previous proofs with a little twist inside.

First, we adjust the definition of shadowy words. A word w is *strongly shadowy* if w is shadowy and for each even position of $\mathbf{w}$ we have that wht and $shdw$ are the most frequent letters among the other labelling $\mathbf{w}$ while for odd positions wht is the most frequent. Note that the words constructed in the previous sections were strongly shadowy because each letter σ appeared only at whites or at shadows.

Exercise 4. There exists an $\text{LTL}_{\mathbf{F}, \mathbf{MFL}}$ formula $\psi_{shadowy}^{MFL}$ defining strongly shadowy words.

Proof. It suffices to revisit Exercise 1 and to modify the formula φ_{odd}^{ex1} stipulating that odd positions are exactly those labelled with $shdw$ (since it is the only formulae employing **Half**). We claim that φ_{odd}^{ex1} can be expressed with

$$\varphi_{odd}^{MFL} := \mathbf{G} [\mathbf{MFL} (wht) \wedge (wht \leftrightarrow \mathbf{MFL} (shdw))]$$

Indeed, take any word $\mathbf{w} \models \varphi_{init}^{ex1} \wedge \varphi_{odd}^{MFL}$. Of course we have $\mathbf{w}, 0 \models wht$ (due to φ_{init}^{ex1}). Moreover, $\mathbf{w}, 1 \models shdw$ holds: otherwise we would get contradiction with $shdw$ not being the most frequent letter in the past of 1. Now assume $p > 1$ and assume that the word $\mathbf{w}_0, \dots, \mathbf{w}_{p-1}$ is strongly shadowy. Consider two cases. If p is odd, then both wht and $shdw$ are the most frequent letters in the past of $p-1$ and $p-1$ is labelled by wht . Then, $shdw$ is not the most frequent letter in the past of p and thus p is labelled by $shdw$ and wht is the most frequent letter in the past of p . If p is even, $p-2$ is labelled by wht and the most frequent letters in the past of $p-2$ are wht and $shdw$, and $p-1$ is labelled by $shdw$. Thus both wht and $shdw$ are the most frequent letters in the past of p and therefore wht is labelled by wht . Thus, $\mathbf{w}_0, \dots, \mathbf{w}_p$ is strongly shadowy. By induction, $\mathbf{w}$ is strongly shadowy. It can be readily checked that every strongly shadowy word satisfies $\psi_{shadowy}^{MFL}$. $\square$

We argue that over the strongly shadowy models, the formulae **Half** σ and **MFL** σ are equivalent.

Lemma 4. For all strongly shadowy words $\mathbf{w} \models \psi_{shadowy}^{MFL}$, all even positions $2i$ and all letters σ we have the equivalence $\mathbf{w}, 2i \models \mathbf{Half} \sigma$ iff $\mathbf{w}, 2i \models \mathbf{MFL} \sigma$.

Proof. If $\mathbf{w}, 2i \models \mathbf{MFL} \sigma$, then $\mathbf{w}, 2i \models \mathbf{MFL} wht$ due to the strong shadowyness of $\mathbf{w}$. Hence $\#_{\sigma}^<(\mathbf{w}, 2i) = \#_{wht}^<(\mathbf{w}, 2i) = \frac{2i}{2}$, implying $\mathbf{w}, 2i \models \mathbf{Half} \sigma$.

Now, assume that $\mathbf{w}, 2i \models \mathbf{Half} \sigma$ holds, so σ appears i times in the past. Since $\mathbf{w}$ is strongly shadowy we know that wht is the most frequent letter. Moreover, wht appears $\frac{2i}{2} = i$ times in the past. Hence, $\mathbf{w}, 2i \models \mathbf{MFL} \sigma$. $\square$

We say that a letter σ is *importunate* in a word $\mathbf{w}$ if σ labels more than half of the positions in some even prefix of $\mathbf{w}$. Notice that strongly shadowy words cannot have importunate letters.

With the above lemma, it is tempting to finish the proof as follows: replace each **Half** (φ) in the formulae from Section 4.1 with **MFL** (p_φ) for some fresh atomic proposition p_φ and require that $\mathbf{G}(\varphi \leftrightarrow p_\varphi)$ holds. A formula obtained from φ in this way will be called a *dehalfication* of φ and will be denoted with $\text{dehalf}(\varphi)$. The next lemma shows that $\text{dehalf}(\cdot)$ preserves satisfaction of certain $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ formulae.

Lemma 5. *Let φ be an $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ formula without nested **Half** operators and without **F** modality, Λ be the set of all formulae λ such that **Half** λ appears in φ and let $\mathbf{w}$ be a word such that $\mathbf{w} \models \psi_{\text{shadowy}}^{\text{MFL}} \wedge \bigwedge_{\lambda \in \Lambda} \mathbf{G}(p_\lambda \leftrightarrow \lambda)$. Then for all even positions $2p$ of $\mathbf{w}$ we have that $\mathbf{w}, 2p \models \text{dehalf}(\varphi)$ implies $\mathbf{w}, 2p \models \varphi$. Moreover, $\mathbf{w} \models \mathbf{G}(\text{whf} \rightarrow \text{dehalf}(\varphi))$ implies $\mathbf{w} \models \mathbf{G}(\text{whf} \rightarrow \varphi)$.*

Proof. The proof goes via structural induction over $\text{LTL}_{\mathbf{F}, \mathbf{Half}}$ formulae without nested **Half** operators and without **F** operators. The only interesting case is when $\varphi = \mathbf{Half} \lambda$, which follows from Lemma 4. $\square$

Note, however, that the above lemma works only one way: it fails when the formula φ is satisfied in more than half of the positions of some prefix, as that would make p_φ importunate leading to unsatisfiability of $\psi_{\text{shadowy}}^{\text{MFL}}$.

6.3 Most-Frequent Letter: the reduction

The next step is to construct a formula defining truly $\Sigma_{\mathcal{A}}$ -shadowy words, which are the crucial part of $\psi_{\text{enc-basics}}^{\text{MFL}}$. To do it, we first need to rewrite a formula $\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}}$, transferring the truth of a letter σ from whites into their shadows. The main ingredient of $\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}}$ is the formula $\varphi(\heartsuit) := \mathbf{G}(\text{whf} \rightarrow \mathbf{Half}([\text{whf} \wedge \sigma] \vee [\text{shdw} \wedge \neg \tilde{\sigma}]))$, which we replace with $\text{dehalf}(\varphi(\heartsuit))$. We call the obtained formula $(\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}})^{\text{MFL}}$ and show its correctness below.

First, by Lemma 5 we know that every model of $(\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}})^{\text{MFL}}$ is also a model of $\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}}$. Then, the models of $\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}}$ can be made strongly shadowy, so dehalfication of $\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}}$ is satisfiability-preserving.

Lemma 6. *Let p_φ be a fresh letter for $\varphi := [\text{whf} \wedge \sigma] \vee [\text{shdw} \wedge \neg \tilde{\sigma}]$. Take $\mathbf{w}$, a strongly shadowy word satisfying $\mathbf{w} \models \varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}}$ without any occurrences of p_φ . Then $\mathbf{w}'$, the word obtained by labelling with p_φ all the positions of $\mathbf{w}$ satisfying φ , is strongly shadowy.*

Hence, we obtain the correctness of $(\varphi_{\sigma \rightsquigarrow \tilde{\sigma}}^{\text{trans}})^{\text{MFL}}$. By applying the same strategy to other conjuncts of $\psi_{\text{enc-basics}}$ and Fact 5, we obtain $\psi_{\text{enc-basics}}^{\text{MFL}}$ satisfying:

Corollary 3. *The function run (taking as input the words satisfying $\psi_{\text{enc-basics}}^{\text{MFL}}$) is uniquely defined and returns words satisfying P1 and P2. Moreover the formulae $\psi_{\text{enc-basics}}^{\text{MFL}}$ and $\psi_{\text{enc-basics}}$ are equi-satisfiable.*

Towards completing the undecidability proof we need to prepare the rewritings of the formulae ψ_{P3} and ψ_{P4} . For ψ_{P3} we proceed similarly to the previous case. We know that the models of $\psi_{enc-basics}^{MFL} \wedge \text{dehalf}(\psi_{P3})$ satisfy P3 (due to Lemma 5 they satisfy ψ_{P3} and hence, by Lemma 2, also P3). To observe the existence of such models, we show again that the satisfiability of ψ_{P3} is preserved by dehalfication.

Lemma 7. *Let p_q be a fresh letter for $\varphi_q := [wht \wedge from_q] \vee [shdw \wedge \neg \widetilde{to_q}]$ indexed over $q \in Q \setminus \{q_0\}$. Take $\mathfrak{w}$, a strongly shadowy word satisfying $\mathfrak{w} \models \psi_{enc-basics}^{MFL} \wedge \psi_{P3}$ without any occurrences of p_q . Then $\mathfrak{w}'$, the word obtained by labelling with p_q all the positions of $\mathfrak{w}$ satisfying φ_q , is strongly shadowy.*

From Lemma 2, Lemma 7 and Lemma 5 we immediately conclude:

Corollary 4. *If $\mathfrak{w}$ satisfies $\psi_{enc-basics}^{MFL} \wedge \text{dehalf}(\psi_{P3})$, then $\text{run}(\mathfrak{w})$ satisfies P1–P3. Moreover the formulae $\psi_{enc-basics}^{MFL} \wedge \text{dehalf}(\psi_{P3})$ and $\psi_{enc-basics} \wedge \psi_{P3}$ are equi-satisfiable.*

The last formula to rewrite is ψ_{P4} . We focus only on its first part, speaking about the first counter, i.e.

$\mathbf{G}(fVal_0 \rightarrow \mathbf{Half}([wht \wedge fOP_{+1}] \vee [shdw \wedge \neg \widetilde{fOP_{-1}}]) \wedge \mathbf{G}(wht \rightarrow (fVal_0 \leftrightarrow \neg fVal_+)))$
 Note that this time we cannot simply dehalfise this formula: the letter responsible for the inner part of **Half** would necessarily be importunate – consider an initial fragment of a run of $\mathcal{A}$ in which $\mathcal{A}$ increments its first counter without decrementing it. Fortunately, we cannot say the same when the machine decrements the counter and hence, it suffices to express the equivalent (due to even length of shadowy models) statement ψ'_{P4} as follows: $\mathbf{G}(fVal_0 \rightarrow \mathbf{Half}(\neg([wht \wedge fOP_{+1}] \vee [shdw \wedge \neg \widetilde{fOP_{-1}}]) \wedge \mathbf{G}(wht \rightarrow (fVal_0 \leftrightarrow \neg fVal_+)))$.

As we did before, we show that dehalfication of ψ'_{P4} preserves satisfiability:

Lemma 8. *Let p_φ be a fresh letter for $\varphi := \neg([wht \wedge fOP_{+1}] \vee [shdw \wedge \neg \widetilde{fOP_{-1}}])$. Take $\mathfrak{w}$, a strongly shadowy word satisfying $\mathfrak{w} \models \psi_{enc-basics}^{MFL} \wedge \text{dehalf}(\psi_{P3}) \wedge \psi'_{P4}$ without any occurrences of p_φ . Then $\mathfrak{w}'$, the word obtained by labelling with p_φ all the positions of $\mathfrak{w}$ satisfying φ , is strongly shadowy.*

Finally, let $(\psi_{\mathcal{A}}^q)^{MFL} := \psi_{enc-basics}^{MFL} \wedge \text{dehalf}(\psi_{P3}) \wedge \text{dehalf}(\psi_{P4}) \wedge \mathbf{F} to_q$. From Lemma 3, Lemma 8 and Lemma 5 we immediately conclude:

Corollary 5. *If $\mathfrak{w}$ satisfies $(\psi_{\mathcal{A}}^q)^{MFL}$ then it satisfies P1–P4. Moreover the formulae $(\psi_{\mathcal{A}}^q)^{MFL}$ and $\psi_{\mathcal{A}}^q$ are equi-satisfiable.*

Thus, by Theorem 1 and the above corollary, we obtain the undecidability of $\text{LTL}_{\mathbf{F}, \mathbf{MFL}}$. Undecidability of the model-checking problem is concluded by virtually the same argument as in Section 6.1. Hence:

Theorem 3. *The model-checking and the satisfiability problems for $\text{LTL}_{\mathbf{F}, \mathbf{MFL}}$ are undecidable.*

7 Decidable variants

We have shown that $\text{LTL}_{\mathbf{F}}$ with frequency operators lead to undecidability. Without the operators that can express $\mathbf{F}$ (e.g. $\mathbf{F}$, $\mathbf{G}$ or $\mathbf{U}$), the decision problems become NP-complete. Below we assume the standard semantics of LTL operator $\mathbf{X}$, i.e. $\mathbf{w}, i \models \mathbf{X} \varphi$ iff $i+1 < |\mathbf{w}|$ and $\mathbf{w}, i+1 \models \varphi$.

Theorem 4. *Model-checking and satisfiability problems for $\text{LTL}_{\mathbf{X}, \mathbf{MFL}, \mathbf{PM}}$ are NP-complete.*

The complexity of $\text{LTL}_{\mathbf{X}, \mathbf{MFL}, \mathbf{PM}}$ is so low because the truth of the formula depends only on some initial fragment of a trace. This is a big restriction of the expressive power. Thus, we consider a different approach motivated by [7].

In the new setting, we allow to use arbitrary LTL formulae as well as percentage operators as long as they are not mixed with $\mathbf{G}$. We introduce a logic $\text{LTL}^{\%}$, which extends the classical LTL [29] with the percentage operators of the form $\mathbf{P}_{\bowtie k\%} \varphi$ for any $\bowtie \in \{\leq, <, =, >, \geq\}$, $k \in \mathbb{N}$ and $\varphi \in \text{LTL}$. By way of example, the formula $\mathbf{P}_{<20\%}(a)$ is true at a position p if less than 20% of positions before p satisfy a . The past majority operator is a special case of the percentage operator: $\mathbf{PM} \equiv \mathbf{P}_{\geq 50\%}$. Formally:

$$\mathbf{w}, i \models \mathbf{P}_{\bowtie k\%} \varphi \text{ if } |\{j < i : \mathbf{w}, j \models \varphi\}| \bowtie \frac{k}{100} i$$

To avoid undecidability, the percentage operators cannot appear under negation or be nested. Therefore, the syntax of $\text{LTL}^{\%}$ is defined with the grammar $\varphi, \varphi' ::= \psi_{\text{LTL}} \mid \varphi \vee \varphi' \mid \varphi \wedge \varphi' \mid \mathbf{F}(\psi_{\text{LTL}} \wedge \mathbf{P}_{\bowtie k\%} \psi'_{\text{LTL}})$, where $\psi_{\text{LTL}}, \psi'_{\text{LTL}}$ are (full) LTL formulae.

The main tool used in the decidability proof is the Parikh Automata [21]. A Parikh automaton $\mathcal{P} = (\mathcal{A}, \mathcal{E})$ over the alphabet Σ is composed of a finite-state automaton $\mathcal{A}$ accepting words from Σ^* and a semi-linear set $\mathcal{E}$ given as a system of linear inequalities with integer coefficients, where the variables are x_a for $a \in \Sigma$. We say that $\mathcal{P}$ accepts a word $\mathbf{w}$ if $\mathcal{A}$ accepts $\mathbf{w}$ and the mapping assigning to each variable x_a from $\mathcal{E}$ the total number of positions of $\mathbf{w}$ carrying the letter a , is a solution to $\mathcal{E}$. Checking non-emptiness of the language of $\mathcal{P}$ can be done in NP [17]. Our main decidability results is obtained by constructing an appropriate Parikh automaton recognising the models of an input $\text{LTL}^{\%}$ formula.

Theorem 5. *Model-checking and satisfiability problems for $\text{LTL}^{\%}$ are decidable.*

Proof. Let $\varphi \in \text{LTL}^{\%}$. By turning φ into a DNF, we can focus on checking satisfiability of some of its conjuncts. Hence, w.l.o.g. we assume that $\varphi = \varphi_0 \wedge \bigwedge_{i=1}^n \varphi_i$, where φ_0 is in LTL and all φ_i have the form $\mathbf{F}(\psi_{\text{LTL}}^{i,1} \wedge \mathbf{P}_{\bowtie k_i\%} \psi_{\text{LTL}}^{i,2})$ for some LTL formulae $\psi_{\text{LTL}}^{i,1}$ and $\psi_{\text{LTL}}^{i,2}$. Observe that a word $\mathbf{w}$ is a model of φ iff it satisfies φ_0 and for each conjunct φ_i we can pick a witness position p_i from $\mathbf{w}$ such that $\mathbf{w}, p_i \models \psi_{\text{LTL}}^{i,1} \wedge \mathbf{P}_{\bowtie k_i\%} \psi_{\text{LTL}}^{i,2}$. Moreover, the percentage constraints inside such formulae speak only about the prefix $\mathbf{w}_{<p_i}$. Thus, knowing the position p_i and the number of positions before p_i satisfying $\psi_{\text{LTL}}^{i,2}$, the percentage constraint inside φ_i can be imposed globally rather than locally. It suggests the use of Parikh

automata: the LTL part of φ can be checked by the appropriate automaton $\mathcal{A}$ (due to the correspondence that for an LTL formula over finite words one can build a finite-state automaton recognising the models of such a formula [19]) and the global constraints, speaking about the satisfaction of percentage operators, can be ensured with a set of linear inequalities $\mathcal{E}$.

Our plan is as follows: we decorate the intended models $\mathbf{w}$ with additional information on witnesses, such that the witness position p_i for φ_i will be labelled by w_i (and there will be a unique such position in a model), all positions before p_i will be labelled by b_i and, among them, we distinguish with a letter s_i some special positions, *i.e.* those satisfying $\psi_{\text{LTL}}^{i,2}$. More formally, for each φ_i we produce an LTL formula φ'_i according to the following rules:

- there is a unique position p_i such that $\mathbf{w}, p_i \models w_i$ (selecting a witness for φ_i),
- for all $j < p_i$ we have $\mathbf{w}, j \models b_i$ (the positions before p_i are labelled with b_i),
- $\mathbf{w} \models \mathbf{G}(s_i \rightarrow [b_i \wedge \psi_{\text{LTL}}^{i,2}])$ (distribution of the special positions among b_i) and
- $\mathbf{w}, p_i \models \psi_{\text{LTL}}^{i,1}$ (a precondition for φ_i).

Let $\varphi' := \varphi_0 \wedge \bigwedge_{i=1}^n \varphi'_i \wedge \bigwedge_{i=1}^n \mathbf{F}(p_i \wedge \mathbf{P}_{\bowtie k_i \% s_i})$. Note that $\mathbf{w} \models \varphi'$ implies $\mathbf{w} \models \varphi$. Moreover, any model $\mathbf{w} \models \varphi$ can be labelled with letters b_i, s_i, w_i such that the decorated word satisfies φ' . Let $\varphi'' := \varphi_0 \wedge \bigwedge_{i=1}^n \varphi'_i$ and let $\mathcal{E}$ be the system of n inequalities with $\mathcal{E}_i = 100 \cdot x_{b_i} \bowtie k_i \cdot x_{s_i}$. Now observe that any model of φ' satisfies $\mathcal{E}$ (*i.e.* the value assigned to x_a is the total number of positions labelled with a), due to the satisfaction of counting operators, and vice versa: every word $\mathbf{w} \models \varphi''$ satisfying $\mathcal{E}$ is a model of φ'' . It gives us a sufficient characterisation of models of φ . Let $\mathcal{A}$ be a finite automaton recognising the models of φ'' , then a Parikh automaton $\mathcal{P} = (\mathcal{A}, \mathcal{E})$, as we already discussed, is non-empty if and only if φ has a model. Since checking non-emptiness of $\mathcal{P}$ is decidable, we can conclude that $\text{LTL}^\%$ is decidable. $\square$

A rough complexity analysis yields an NEXPTIME upper bound on the problem: the automaton $\mathcal{P}$ that we constructed is exponential in φ (translating φ to DNF does not increase the complexity since we only guess one conjunct, which is of polynomial size in φ). Moreover, checking non-emptiness can be done non-deterministically in time polynomial in the size of the automaton. The NEXPTIME bound is not optimal: we conjecture that the problem is PSPACE-complete. We believe that by employing techniques similar to [7], one can construct $\mathcal{P}$ and check its non-emptiness on the fly, which should result in the PSPACE upper bound.

For the model-checking problem, we observe that determining whether some trace of a Kripke structure $\mathcal{K} = (S, I, R, l)$ satisfies φ is equivalent to checking the satisfiability of formula $\varphi_{\mathcal{K}} \wedge \varphi$, where $\varphi_{\mathcal{K}}$ is a formula describing all the traces of $\mathcal{K}$. Such a formula can be constructed in a standard manner. For simplicity, we treat S as a set of auxiliary letters, and consider the conjunction of (1) $\bigvee_{s \in I} s$, (2) $\mathbf{G}(\mathbf{X} \top \rightarrow \bigvee_{(s,s') \in R} (s \wedge \mathbf{X} s'))$ and (3) $\bigwedge_{s \in S} \mathbf{G}(s \rightarrow \bigwedge_{p \in \ell(s)} p)$, expressing that the trace starts with an initial state, consecutive positions describe consecutive states and that the trace is labelled by the appropriate letters. Thus, the model-checking problem can be reduced in polynomial time to the satisfiability problem.

8 Two-Variable First-Order Logic with Majority

The *Two-Variable First-Order Logic on words* ($\text{FO}^2[<]$) is a robust fragment of First-Order Logic FO interpreted on finite words. It involves quantification over variables x and y (ranging over the words' positions) and it admits a linear order predicate $<$ (interpreted as a natural order on positions) and the equality predicate $=$. Henceforth we assume the usual semantics of $\text{FO}^2[<]$ (cf. [16]).

In this section, we investigate the logic $\text{FO}_M^2[<]$, namely the extension of $\text{FO}^2[<]$ with the so-called *Majority quantifier* M . Such quantifier was intensively studied due to its close connection with circuit complexity and algebra, see e.g. [22, 5, 6]. Intuitively, the formula $Mx.\varphi$ specifies that at least half of all the positions in a model, after substituting x with them, satisfy φ . Formally $\mathfrak{w} \models Mx.\varphi$ holds, if and only if $\frac{|\mathfrak{w}|}{2} \leq |\{p \mid \mathfrak{w}, p \models \varphi[x/p]\}|$. We stress that the formula $Mx.\varphi$ may contain free occurrences of the variable y .

Note that the Majority quantifier shares similarities to the **PM** operator, but in contrast to **PM**, the M quantifier counts *globally*. We take advantage of such similarities and by reusing the technique developed in the previous sections, we show that the satisfiability problem for $\text{FO}_M^2[<]$ is also undecidable. We stress that our result significantly sharpens an existing undecidability result for FO with Majority from [23] (since in our case the number of variables is limited) as well as for $\text{FO}^2[<, succ]$ with Presburger Arithmetics from [25] (since our counting mechanism is limited and the successor relation *succ* is disallowed).

Proof plan There are three possible approaches to proving the undecidability of $\text{FO}_M^2[<]$. The first one is to reproduce all the results for $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$, which is rather uninspiring. The second one is to define a translation from $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$ to $\text{FO}_M^2[<]$ that produces an equisatisfiable formula. But because of models of odd length, this involves a lot of case study. Here we present a third approach, which, we believe, gives the best insight: we show a translation from $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$ to $\text{FO}_M^2[<]$ that works for $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$ formulae whose all models are shadowy. Since we only use such models in the undecidability proof of $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$, this shows the undecidability of $\text{FO}_M^2[<]$.

Shadowy models We first focus on defining shadowy words in $\text{FO}_M^2[<]$. Before we start, let us introduce a bunch of useful macros in order to simplify the forthcoming formulae. Their names coincide with their intuitive meaning and their semantics.

- $\text{Half}x.\varphi := Mx.\varphi \wedge Mx.\neg\varphi$,
- $\text{first}(x) := \neg\exists y y < x$, $\text{second}(x) := \exists y y < x \wedge \forall y y < x \rightarrow \text{first}(y)$,
- $\text{last}(x) := \neg\exists y y > x$, $\text{sectolast}(x) := \exists y y > x \wedge \forall y y > x \rightarrow \text{last}(y)$

Lemma 9. *There is an $\text{FO}_M^2[<]$ formula $\psi_{\text{shadowy}}^{\text{FO}}$ defining shadowy words.*

Proof. Let $\varphi_{\text{base}}^{\text{lem9}}$ be a formula defining the language of all (non-empty) words, where the letters *wht* and *shdw* label disjoint positions in the way that the first position satisfies *wht* and the total number of *shdw* and *wht* coincide. It can be

written, *e.g.* with $\forall x(wht(x) \leftrightarrow \neg shdw(x)) \wedge \exists x(first(x) \wedge wht(x)) \wedge \text{Half}x.wht(x) \wedge \text{Half}x.shdw(x)$. To define shadowy words, it would be sufficient to specify that no neighbouring positions carry the same letter among $\{wht, shdw\}$. This can be done with, rather complicated at the first glance, formulae:

$$\varphi_{wht.wht}^{forbid}(x) := wht(x) \rightarrow \text{Half}y. ([y < x \wedge wht(y)] \vee [x < y \wedge shdw(y)]),$$

$$\varphi_{shdw.shdw}^{forbid}(x) := shdw(x) \rightarrow \text{Half}y. ([y < x \vee x = y] \wedge shdw(y)) \vee [x < y \wedge wht(y)].$$

Finally, let $\psi_{shadowy}^{FO} := \varphi_{base}^{lem9} \wedge \forall x. (\varphi_{wht.wht}^{forbid}(x) \wedge \varphi_{shdw.shdw}^{forbid}(x))$.

Showing that shadowness implies the satisfaction of $\psi_{shadowy}^{FO}$ can be done by routine induction. For the opposite direction, take $\mathfrak{w} \models \psi_{shadowy}^{FO}$. Since $\mathfrak{w} \models \varphi_{base}^{lem9}$ the only possibility for $\mathfrak{w}$ to not be shadowy is to have two consecutive positions $p, p+1$ carrying the same letter. W.l.o.g assume they are both white. Let w be the number of white positions to the left of p and let s be the number of shadows to the right of p . By applying $\varphi_{wht.wht}^{forbid}$ to p we infer that $w + s = \frac{1}{2}|\mathfrak{w}|$. On the other hand, by applying $\varphi_{wht.wht}^{forbid}$ to $p+1$ it follows that $(w+1)+s = \frac{1}{2}|\mathfrak{w}|$, which contradicts the previous equation. Hence, $\mathfrak{w}$ is shadowy. $\square$

Translation It is a classical result from [16] that $\text{FO}^2[<]$ can express $\text{LTL}_{\mathbf{F}}$. We define a translation $\text{tr}_v(\varphi)$ from $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$ to $\text{FO}_{\mathbf{M}}^2[<]$, parametrised by a variable v (where v is either x or y and $\bar{v}$ denotes the different variable from v), inductively. We write $v \leq \bar{v}$ rather than $v < \bar{v} \vee v = \bar{v}$ for simplicity. For $\text{LTL}_{\mathbf{F}}$ cases, we follow [16]: $\text{tr}_v(a) := a(v)$, for a fresh unary predicate a for each $a \in \text{AP}$, $\text{tr}_v(\neg\varphi) := \neg\text{tr}_v(\varphi)$, $\text{tr}_v(\varphi \wedge \varphi') := \text{tr}_v(\varphi) \wedge \text{tr}_v(\varphi')$, $\text{tr}_v(\mathbf{F}\varphi) := \exists \bar{v} (v \leq \bar{v}) \wedge \text{tr}_{\bar{v}}(\varphi)$. For $\mathbf{PM}$, we propose $\text{tr}_v(\mathbf{PM}\varphi) := \text{M}\bar{v}((\bar{v} < v \wedge \text{tr}_{\bar{v}}(\varphi)) \vee (\bar{v} \geq v \wedge wht(\bar{v})))$. Finally, for a given $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$ formula φ , let $\text{tr}(\varphi)$ stand for $\psi_{shadowy}^{FO} \wedge \exists x.(first(x) \wedge \text{tr}_x(\varphi))$.

The following lemma shows the correctness of the presented translation.

Lemma 10. *An $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$ formula φ has a shadowy model iff $\text{tr}(\varphi)$ has a model.*

Since the formulae used in our undecidability proof for $\text{LTL}_{\mathbf{F}, \mathbf{PM}}$ have only shadowy models, by Lemma 10 we conclude that $\text{FO}_{\mathbf{M}}^2[<]$ is also undecidable.

Theorem 6. *The satisfiability problem for $\text{FO}_{\mathbf{M}}^2[<]$ is undecidable.*

9 Conclusions

We have provided a simple proof showing that adding different percentage operators to $\text{LTL}_{\mathbf{F}}$ yields undecidability. We showed that our technique can be applied to an extension of first-order logic on words, and we hope that our work will turn useful in showing undecidability for other extensions of temporal logics. Decidability results for logics with percentage operators in restricted contexts were also provided.

Acknowledgements

Bartosz Bednarczyk was supported by the Polish Ministry of Science and Higher Education program “Diamantowy Grant” no. DI2017 006447. Jakub Michaliszyn was supported by NCN grant no. 2017/27/B/ST6/00299.

References

1. Jorge A. Baier, Sheila A. McIlraith: Planning with First-Order Temporally Extended Goals using Heuristic Search. In: AAAI (2006)
2. Shaull Almagor, Udi Boker, Orna Kupferman: Discounting in LTL. In: TACAS (2014). https://doi.org/10.1007/978-3-642-54862-8_37
3. Rajeev Alur, Thomas A. Henzinger: Real-Time Logics: Complexity and Expressiveness. Inf. Comput. (1993)
4. Bednarczyk, B., Michaliszyn, J.: "most of" leads to undecidability: Failure of adding frequencies to LTL <https://arxiv.org/abs/2007.01233>
5. Christoph Behle, Andreas Krebs: Regular Languages in MAJ[>] with three variables. ECCC (2011)
6. Christoph Behle, Andreas Krebs, Stephanie Reifferscheid: Regular Languages Definable by Majority Quantifiers with Two Variables. In: DLT (2009). https://doi.org/10.1007/978-3-642-02737-6_7
7. Udi Boker, Krishnendu Chatterjee, Thomas A. Henzinger, Orna Kupferman: Temporal Specifications with Accumulative Values. ACM Trans. Comput. Log. (2014). <https://doi.org/10.1145/2629686>
8. Benedikt Bollig, Normann Decker, Martin Leucker: Frequency Linear-time Temporal Logic. In: TASE (2012). <https://doi.org/10.1109/TASE.2012.43>
9. Patricia Bouyer, Nicolas Markey, Raj Mohan Matteplackel: Averaging in LTL. In: CONCUR (2014). https://doi.org/10.1007/978-3-662-44584-6_19
10. Büchi, J.R.: Weak Second-Order Arithmetic and Finite Automata. Mathematical Logic Quarterly (1960). <https://doi.org/10.1002/malq.19600060105>
11. Cadilhac, M., Finkel, A., McKenzie, P.: Affine Parikh automata. RAIRO Theor. Informatics Appl. (2012). <https://doi.org/10.1051/ita/2012013>
12. Witold Charatonik, Piotr Witkowski: Two-variable Logic with Counting and a Linear Order. Logical Methods in Computer Science (2016). [https://doi.org/10.2168/LMCS-12\(2:8\)2016](https://doi.org/10.2168/LMCS-12(2:8)2016)
13. Normann Decker, Peter Habermehl, Martin Leucker, Arnaud Sangnier, Daniel Thoma: Model-Checking Counting Temporal Logics on Flat Structures. In: CONCUR (2017). <https://doi.org/10.4230/LIPIcs.CONCUR.2017.29>
14. Stéphane Demri: LTL over integer periodicity constraints. Theor. Comput. Sci. (2006). <https://doi.org/10.1016/j.tcs.2006.02.019>
15. Demri, S., Goranko, V., Lange, M.: Temporal Logics in Computer Science: Finite-State Systems. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press (2016). <https://doi.org/10.1017/CBO9781139236119>
16. Kousha Etessami, Moshe Y. Vardi, Thomas Wilke: First-Order Logic with Two Variables and Unary Temporal Logic. Inf. Comput. (2002). <https://doi.org/10.1006/inco.2001.2953>
17. Diego Figueira, Leonid Libkin: Path Logics for Querying Graphs: Combining Expressiveness and Efficiency. In: LICS (2015). <https://doi.org/10.1109/LICS.2015.39>
18. Giuseppe De Giacomo, Moshe Y. Vardi: Linear Temporal Logic and Linear Dynamic Logic on Finite Traces. In: IJCAI (2013)
19. Giuseppe De Giacomo, Moshe Y. Vardi: Synthesis for LTL and LDL on Finite Traces. In: IJCAI (2015)
20. Jochen Hoenicke, Roland Meyer, Ernst-Rüdiger Olderog: Kleene, Rabin, and Scott Are Available. In: CONCUR (2010). https://doi.org/10.1007/978-3-642-15375-4_32
21. Felix Klaedtke, Harald Rueß: Monadic Second-Order Logics with Cardinalities. In: ICALP (2003). https://doi.org/10.1007/3-540-45061-0_54

22. Andreas Krebs: Typed semigroups, majority logic, and threshold circuits. Ph.D. thesis, University of Tübingen, Germany (2008)
23. Klaus-Jörn Lange: Some Results on Majority Quantifiers over Words. In: CCC (2004). <https://doi.org/10.1109/CCC.2004.1313817>
24. François Laroussinie, Antoine Meyer, Eudes Pettonnet: Counting LTL. In: TIME (2010). <https://doi.org/10.1109/TIME.2010.20>
25. Kamal Lodaya, Sreejith, A.V.: Two-Variable First Order Logic with Counting Quantifiers: Complexity Results. In: DLT (2017). https://doi.org/10.1007/978-3-319-62809-7_19
26. Minsky, M.L.: Computation: Finite and Infinite Machines. Prentice-Hall Series in Automatic Computation, Prentice-Hall (1967)
27. Matthias Niewerth: Data definition languages for XML repository management systems. Ph.D. thesis, Technical University of Dortmund, Germany (2016)
28. Joël Ouaknine, James Worrell: On the decidability and complexity of Metric Temporal Logic over finite words. Logical Methods in Computer Science (2007)
29. Amir Pnueli: The Temporal Logic of Programs. In: FOCS (1977). <https://doi.org/10.1109/SFCS.1977.32>

Open Access This chapter is licensed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license and indicate if changes were made.

The images or other third party material in this chapter are included in the chapter’s Creative Commons license, unless indicated otherwise in a credit line to the material. If material is not included in the chapter’s Creative Commons license and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder.