

Complexity of Safety and coSafety Fragments of Linear Temporal Logic

Alessandro Artale¹, Luca Geatti^{2*}, Nicola Gigante¹, Andrea Mazzullo^{1*}, Angelo Montanari²

¹Free University of Bozen-Bolzano

² University of Udine

artale@inf.unibz.it, luca.geatti@uniud.it, gigante@inf.unibz.it, mazzullo@inf.unibz.it, angelo.montanari@uniud.it

Abstract

Linear Temporal Logic (LTL) is the de-facto standard temporal logic for system specification, whose foundational properties have been studied for over five decades. Safety and cosafety properties define notable fragments of LTL, where a prefix of a trace suffices to establish whether a formula is true or not over that trace. In this paper, we study the complexity of the problems of satisfiability, validity, and realizability over infinite and finite traces for the safety and cosafety fragments of LTL. As for satisfiability and validity over infinite traces, we prove that the majority of the fragments have the same complexity as full LTL, that is, they are PSPACE-complete. The picture is radically different for realizability: we find fragments with the same expressive power whose complexity varies from 2EXPTIME-complete (as full LTL) to EXPTIME-complete. Notably, for all cosafety fragments, the complexity of the three problems does not change passing from infinite to finite traces, while for all safety fragments the complexity of satisfiability (resp., realizability) over finite traces drops to NP-complete (resp., Π_2^P -complete).

1 Introduction

Linear Temporal Logic (LTL) is arguably the most renowned temporal logic, with applications in a variety of branches of computer science (Pnueli 1977; Vardi and Wolper 1986, 1994). LTL is usually interpreted over infinite state sequences (or traces); recently, the finite-trace semantics has received attention as well, especially in *artificial intelligence* (De Giacomo and Vardi 2013, 2015; De Giacomo, De Masellis, and Montali 2014; Fionda and Greco 2018; Artale, Mazzullo, and Ozaki 2019).

The satisfiability (resp., validity) problem of LTL consists of deciding whether, given an LTL formula, it is satisfied by at least one state sequence (resp., by all state sequences). It is known that satisfiability and validity of LTL, interpreted over both infinite and finite traces, are PSPACE-complete (Sistla and Clarke 1985; De Giacomo and Vardi 2013). Realizability (Pnueli and Rosner 1989) is more complex than satisfiability: it asks, for a given formula over a set of variables partitioned into controllable and uncontrollable ones, whether there exists a strategy such that, for any value of the uncontrollable variables, chooses the value of the controllable

ones in such a way to satisfy the formula. LTL realizability is 2EXPTIME-complete, on both infinite (Rosner 1992) and finite (De Giacomo and Vardi 2015) traces.

Despite the complexity of these problems, several LTL tools have been developed, including model checkers and translators to automata. However, some applications (such as in *runtime verification*) do not always require the full expressivity of LTL, and would rather benefit instead from computational efficiency. Several fragments considered in the literature address these aspects. Two notable ones are the *safety* and *cosafety* fragments (Sistla 1994): they are a subclass of ω -regular languages where a finite prefix suffices to establish the membership of an infinite word to a language, thus allowing one to reason over finite traces. This is very helpful in practice, *e.g.*, it allows one to avoid Safra’s determinization algorithm (Safra 1988) in favor of the classical subset construction. However, despite their usefulness, a systematic complexity analysis of reasoning in these fragments, over both infinite and finite traces, is missing.

In this paper, we study the complexity of the satisfiability, validity, and realizability problems for safety and cosafety LTL fragments, over both infinite and finite traces. We focus on three cosafety fragments and the dual safety ones, some of which are expressively complete with respect to the set of LTL-definable (co)safety properties.

We first prove that the complexity of the satisfiability and validity problems for the majority of the considered fragments (both safety and cosafety) is the same as full LTL, that is, PSPACE-complete.

As for cosafety fragments, we prove a general theorem that allows us to transfer all complexity results for satisfiability, validity, and realizability from infinite to finite trace semantics. On the contrary, the difference in complexity when passing from infinite to finite traces is not negligible for safety fragments. We prove a small (bounded) model property for all safety regular languages of finite words, which states that if a language is not empty, then it contains a word of length 1. By exploiting this result, we show that the complexity of satisfiability (resp., realizability) of all safety fragments drops to NP-complete (resp., Π_2^P -complete) when considering finite traces.

Finally, we show that some of the fragments, although being expressively equivalent, have different complexities for realizability. In particular, for fragments that use past modal-

*Corresponding author.

ities (or are devoid of *until* and *release* modalities), the complexity turns out to be EXPTIME-complete, in contrast to the 2EXPTIME-completeness of the other fragments.

The paper is organized as follows. In Sec. 2, we provide the necessary background. Sec. 3 contains two general theorems that we will use for establishing the complexities of the fragments under finite trace semantics. Sec. 4 and Sec. 5 study the complexity of the fragments, over both infinite and finite traces, for the satisfiability/validity and realizability problems, respectively. In Sec. 6 we discuss the results, while in Sec. 7 we point out future research directions.

The proof of all Lemmata and Theorems can be found in (Artale et al. 2022).

2 Preliminaries

In this section, we provide the necessary background.

Linear Temporal Logic. Given a set Σ of proposition letters, an LTL+P formula ϕ is generated as follows:

$$\begin{array}{ll} \phi := p \mid \neg p \mid \phi \vee \psi \mid \phi \wedge \psi & \text{Boolean connectives} \\ \mid X\phi \mid \tilde{X}\phi \mid \phi U \psi \mid \phi R \psi & \text{future modalities} \\ \mid Y\phi \mid \tilde{Y}\phi \mid \phi S \psi \mid \phi T \psi & \text{past modalities} \end{array}$$

where $p \in \Sigma$. We use the standard shortcuts for $\top := p \vee \neg p$, $\perp := p \wedge \neg p$ (for some $p \in \Sigma$) and other temporal operators: $F\phi := \top U \phi$, $G\phi := \perp R \phi$, $P\phi := \top S \phi$, and $H\phi := \perp T \phi$. Note that, *w.l.o.g.*, our definition of LTL+P considers formulas already in Negation Normal Form (NNF), that is, negations are applied only to proposition letters.

A *pure future* (resp., *past*) formula is an LTL+P formula without occurrences of past (resp., future) modalities. We denote by LTL (resp., pLTL) the set of pure future (resp., past) formulas. Given a set $S \subseteq \{X, \tilde{X}, U, R, Y, \tilde{Y}, S, T\}$ of temporal operators and $\mathbb{L} \in \{\text{LTL+P}, \text{LTL}, \text{pLTL}\}$, we denote by $\mathbb{L}[S]$ the set of formulas ϕ of \mathbb{L} restricted to operators in S . In the following, we denote by Safety-LTL (resp., coSafety-LTL) the fragment LTL $[\tilde{X}, R]$ (resp., LTL $[X, U]$), also known as the *syntactic (co-)safety fragment of LTL* (Sistla 1985; Chang, Manna, and Pnueli 1992; Zhu et al. 2017). Finally, we denote by G(pLTL) (resp., F(pLTL)) the set of LTL+P formulas of the form $G\alpha$ (resp., $F\alpha$), with $\alpha \in \text{pLTL}$.

Let $\sigma \in (2^\Sigma)^+ \cup (2^\Sigma)^\omega$ be a *state sequence* (or *trace*, or *word*). We define the *length* of σ as $|\sigma| = n$, if $\sigma = \langle \sigma_0, \dots, \sigma_{n-1} \rangle \in (2^\Sigma)^+$; $|\sigma| = \omega$, if $\sigma \in (2^\Sigma)^\omega$. The *satisfaction* of an LTL+P formula ϕ by σ at time $0 \leq i < |\sigma|$, denoted by $\sigma, i \models \phi$, is defined as follows (we omit Booleans):

- $\sigma, i \models p$ iff $p \in \sigma_i$;
- $\sigma, i \models X\phi$ iff $i + 1 < |\sigma|$ and $\sigma, i + 1 \models \phi$;
- $\sigma, i \models \tilde{X}\phi$ iff either $i + 1 = |\sigma|$ or $\sigma, i + 1 \models \phi$;
- $\sigma, i \models Y\phi$ iff $i > 0$ and $\sigma, i - 1 \models \phi$;
- $\sigma, i \models \tilde{Y}\phi$ iff either $i = 0$ or $\sigma, i - 1 \models \phi$;
- $\sigma, i \models \phi_1 U \phi_2$ iff there exists $i \leq j < |\sigma|$ such that $\sigma, j \models \phi_2$, and $\sigma, k \models \phi_1$ for all k , with $i \leq k < j$;
- $\sigma, i \models \phi_1 S \phi_2$ iff there exists $j \leq i$ such that $\sigma, j \models \phi_2$, and $\sigma, k \models \phi_1$ for all k , with $j < k \leq i$;

- $\sigma, i \models \phi_1 R \phi_2$ iff either $\sigma, j \models \phi_2$ for all $i \leq j < |\sigma|$, or there exists $i \leq k < |\sigma|$ such that $\sigma, k \models \phi_1$ and $\sigma, j \models \phi_2$ for all $i \leq j \leq k$;
- $\sigma, i \models \phi_1 T \phi_2$ iff either $\sigma, j \models \phi_2$ for all $0 \leq j \leq i$, or there exists $k \leq i$ such that $\sigma, k \models \phi_1$ and $\sigma, j \models \phi_2$ for all $i \geq j \geq k$.

We say that σ is a *model* of ϕ (written as $\sigma \models \phi$) iff $\sigma, 0 \models \phi$. The *language of infinite* (resp., *finite*) *traces* of ϕ , denoted by $\mathcal{L}(\phi)$, is the set of traces $\sigma \in (2^\Sigma)^\omega$ (resp., $\sigma \in (2^\Sigma)^+$) such that $\sigma \models \phi$. We say that two formulas $\phi, \psi \in \text{LTL+P}$ are *equivalent on infinite* (resp., *finite*) *traces*, written $\phi \equiv_I \psi$ (resp., $\phi \equiv_F \psi$), when, for all $\sigma \in (2^\Sigma)^\omega$ (resp., $\sigma \in (2^\Sigma)^+$), it holds that σ is a model of ϕ if and only if σ is a model of ψ . We simply write \equiv when it is clear from the context which one between \equiv_F and \equiv_I has to be used.

If ϕ belongs to pLTL (*i.e.*, pure past fragment of LTL+P), then we interpret ϕ only on *finite* state sequences and we say that $\sigma \in (2^\Sigma)^+$ is a model of ϕ if and only if $\sigma, |\sigma| - 1 \models \phi$, that is, each ϕ in pLTL is interpreted at the *last* state of a finite state sequence.

Safety and Cosafety Fragments of LTL. We recall the definition of safety and cosafety ω -regular languages. Let A be a finite alphabet. For any $\sigma \in A^+ \cup A^\omega$ and any $i < |\sigma|$, we denote by $\sigma_{[0,i]}$ the prefix of σ from 0 to i .

Definition 1 (Co-safety language (Kupferman and Vardi 2001; Thomas 1988)). *Let $\mathcal{L} \subseteq A^\omega$ (resp., $\mathcal{L} \subseteq A^+$). We say that \mathcal{L} is a co-safety language of infinite (resp., finite) words if and only if for all $\sigma \in A^\omega$ (resp. $\sigma \in A^+$), it holds that if $\sigma \in \mathcal{L}$, then there exists $i \in \mathbb{N}$ (resp., $i < |\sigma|$) such that $\sigma_{[0,i]} \cdot \sigma' \in \mathcal{L}$, for all $\sigma' \in A^\omega$ (resp., $\sigma' \in A^+$).*

Definition 2 (Safety language). *A language \mathcal{L} is a safety language iff its complement $\bar{\mathcal{L}}$ is a cosafety language.*

Let $\mathbb{L} \subseteq \text{LTL+P}$. We say that \mathbb{L} is a safety (resp., cosafety) fragment of LTL+P iff $\mathcal{L}(\phi)$ is a safety (resp., cosafety) language, for any $\phi \in \mathbb{L}$. The following result establishes a connection between the semantic and the syntactic (co)safety fragment of LTL+P.

Proposition 1 (Chang, Manna, and Pnueli (1992); Thomas (1988); Cimatti et al. (2022)). *Let ϕ be a formula of LTL and let $\mathcal{L}(\phi)$ be the language of ϕ over infinite or over finite traces. The following sentences are equivalent:*

- $\mathcal{L}(\phi)$ is a safety (resp., co-safety) language;
- there exists a formula ϕ' in G(pLTL) (resp., F(pLTL)) such that $\phi' \equiv \phi$;
- there exists a formula ϕ'' in Safety-LTL (resp., coSafety-LTL) such that $\phi'' \equiv \phi$.

Satisfiability and Validity. We say that an LTL+P formula ϕ is *satisfiable on infinite* (resp., *finite*) *traces* if there exists a trace $\sigma \in (2^\Sigma)^\omega$ (resp., $\sigma \in (2^\Sigma)^+$) such that σ is a model of ϕ . We say that ϕ is *valid on infinite* (resp., *finite*) *traces* if, for every trace $\sigma \in (2^\Sigma)^\omega$ (resp., $\sigma \in (2^\Sigma)^+$), we have that σ is a model of ϕ .

Given a set of formulas \mathbb{L} , the *satisfiability* (resp., *validity*) *problem for \mathbb{L}* on finite or infinite traces, respectively, is the problem of establishing, given a formula $\phi \in \mathbb{L}$, whether ϕ

is satisfiable (resp., valid) on infinite or finite traces, respectively. We recall some results from the literature on the complexity of the satisfiability and validity problems of (fragments of) LTL on infinite and finite traces.

Proposition 2 (Sistla and Clarke (1985); De Giacomo and Vardi (2013)). *The satisfiability problems for LTL and LTL+P (resp., LTL[X, F]) on infinite and on finite traces is PSPACE-complete (resp., NP-complete).*

Proposition 3 (Cf. e.g. Gabbay et al. (2003), Section 1.6). *Let \mathbb{L} and \mathbb{L}' be two sets of formulas such that $\phi \in \mathbb{L}$ iff the transformation into NNF of $\neg\phi \in \mathbb{L}'$, and let \mathbb{C} be a complexity class. It holds that the satisfiability problem for \mathbb{L} is \mathbb{C} -complete iff the validity problem for \mathbb{L}' is $\text{co}\mathbb{C}$ -complete.*

From Propositions 2 and 3, one can prove the following result on the complexity of the validity problem for LTL, LTL+P and LTL[X, G].

Proposition 4. *The validity problem for LTL and LTL+P (resp., LTL[X, G]) on infinite and on finite traces is PSPACE-complete (resp., coNP -complete).*

Realizability. We define the realizability problem for temporal logic formulas as a two-player game between Controller, whose aim is to satisfy the formula, and Environment, who tries to violate it. In this setting, the notion of strategy plays a crucial role.

Definition 3 (Strategy). *Let $\Sigma = \mathcal{C} \cup \mathcal{U}$ be a set of variables partitioned into controllable \mathcal{C} and uncontrollable \mathcal{U} ones. A strategy for Controller is a function $s : (2^{\mathcal{U}})^+ \rightarrow 2^{\mathcal{C}}$ that, for any finite sequence $U = \langle U_0, \dots, U_n \rangle$ of choices by Environment, determines the choice $C_n = s(U)$ of Controller.*

Let $s : (2^{\mathcal{U}})^+ \rightarrow 2^{\mathcal{C}}$ be a strategy and let $U = \langle U_0, U_1, \dots \rangle \in (2^{\mathcal{U}})^\omega$ be an infinite sequence of choices by Environment. We denote by $\text{res}(s, U) = \langle U_0 \cup s(\langle U_0 \rangle), U_1 \cup s(\langle U_0, U_1 \rangle), \dots \rangle$ the state sequence resulting from reacting to U according to s . The realizability problem can be defined as follows.

Definition 4 (Realizability). *Let ϕ be an LTL+P formula over the alphabet $\Sigma = \mathcal{C} \cup \mathcal{U}$, with $\mathcal{C} \cap \mathcal{U} = \emptyset$. We say that ϕ is realizable over infinite (resp., finite) traces if and only if there exists a strategy $s : (2^{\mathcal{U}})^+ \rightarrow 2^{\mathcal{C}}$ such that, for any infinite sequence $U = \langle U_0, U_1, \dots \rangle$ in $(2^{\mathcal{U}})^\omega$, it holds that $\text{res}(s, U) \models \phi$ (resp., there exists $k \in \mathbb{N}$ such that the prefix of $\text{res}(s, U)$ from 0 to k is a model of ϕ).*

Given a set of formulas \mathbb{L} , the realizability problem of \mathbb{L} is the problem of establishing, given a formula $\phi \in \mathbb{L}$, whether ϕ is realizable. We recall some results in the literature on the complexity of the realizability problem of (fragments of) LTL and LTL+P over infinite and finite traces.

Proposition 5 ((Pnueli and Rosner 1989; Rosner 1992; De Giacomo and Vardi 2015)). *Realizability for LTL and LTL+P over infinite and over finite traces is 2EXPTIME-complete.*

3 General Results on Finite Traces

In this section, we provide some theorems that we will use in the following to determine the complexities of (co)safety fragments interpreted over finite traces.

We begin with the definition of *suffix independence* for a logic, which requires infinite models of its formulas to coincide with the concatenation of finite models with arbitrary infinite traces.

Definition 5. *Let \mathbb{L} be a fragment of LTL+P. We say that \mathbb{L} is suffix independent iff, for any $\phi \in \mathbb{L}$ over the alphabet Σ , $\mathcal{L}(\phi) = \mathcal{L}(\phi)_F \cdot (2^\Sigma)^\omega$, where $\mathcal{L}(\phi)$ (resp. $\mathcal{L}(\phi)_F$) is the language of ϕ over infinite (resp. finite) traces.*

For suffix independent logics, we prove the following equisatisfiability result: if a formula is satisfiable on infinite traces, it is satisfiable also over finite traces, and vice versa.

Theorem 1. *Let \mathbb{L} be a fragment of LTL+P that is suffix independent. For any $\phi \in \mathbb{L}$, it holds that:*

$$\mathcal{L}(\phi) \neq \emptyset \Leftrightarrow \mathcal{L}(\phi)_F \neq \emptyset$$

where $\mathcal{L}(\phi)$ (resp. $\mathcal{L}(\phi)_F$) is the language of ϕ over infinite (resp. finite) traces.

The second theorem of this section is a small (bounded) model property for all safety languages of finite words, which proves that if any of these languages is not empty, then there is at least a word of length 1 in the language.

Theorem 2. *Let $\mathcal{L} \subseteq (2^\Sigma)^+$ be a safety language of finite traces. If $\mathcal{L} \neq \emptyset$, then there exists a word $\langle \sigma_0 \rangle$ of length 1 such that $\langle \sigma_0 \rangle \in \mathcal{L}$.*

Theorem 2 will let us prove that the complexity of the satisfiability and realizability problems of safety fragments significantly decreases when passing from infinite to finite words. The next result proves a stronger property of the safety fragments of LTL.

Theorem 3. *Let $\mathbb{L} \in \{\text{LTL}[\tilde{X}, G], \text{Safety-LTL}, G(\text{pLTL})\}$. For any $\phi \in \mathbb{L}$ and any $\sigma = \langle \sigma_0, \dots, \sigma_n \rangle \in (2^\Sigma)^+$ (for some $n \in \mathbb{N}$), if $\sigma \models \phi$ then $\langle \sigma_0 \rangle \models \phi$.*

4 Complexity of Satisfiability and Validity

In this section, we study the complexity of the satisfiability and validity problems for safety and cosafety fragments of LTL+P on both infinite and finite traces. In particular, here and in the rest of the paper, we will focus on the cosafety fragments coSafety-LTL , $F(\text{pLTL})$, and $\text{LTL}[X, F]$, and the dual safety fragments Safety-LTL , $G(\text{pLTL})$, and $\text{LTL}[\tilde{X}, G]$. On infinite traces, we show the following results.

Theorem 4. *The satisfiability and validity problems on infinite traces are PSPACE-complete for:*

1. coSafety-LTL , $F(\text{pLTL})$;
2. $\text{LTL}[\tilde{X}, G]$, Safety-LTL , $G(\text{pLTL})$.

Moreover, we prove the following results on finite traces.

Theorem 5. *The satisfiability problem on finite traces is:*

1. PSPACE-complete for coSafety-LTL , $F(\text{pLTL})$;
2. NP-complete for $\text{LTL}[\tilde{X}, G]$, Safety-LTL , $G(\text{pLTL})$.

The validity problem on finite traces is coNP -complete for $\text{LTL}[X, F]$, coSafety-LTL , $F(\text{pLTL})$.

Logics	Problems					
	satisfiability		validity		realizability	
	infinite	finite	infinite	finite	infinite	finite
coSafety-LTL	PSPACE-c	PSPACE-c	PSPACE-c	coNP-c	2EXPTIME-c	2EXPTIME-c
F(pLTL)	PSPACE-c	PSPACE-c	PSPACE-c	coNP-c	EXPTIME-c	EXPTIME-c
LTL[X, F]	NP-c	NP-c	PSPACE-c	coNP-c	EXPTIME-c	EXPTIME-c

Table 1: Table with complexity results on the cosafety fragments of LTL (the results in grey are proved in this paper).

Logics	Problems					
	satisfiability		validity		realizability	
	infinite	finite	infinite	finite	infinite	finite
Safety-LTL	PSPACE-c	NP-c	PSPACE-c	PSPACE-c	2EXPTIME-c	Π_2^P -c
G(pLTL)	PSPACE-c	NP-c	PSPACE-c	PSPACE-c	EXPTIME-c	Π_2^P -c
LTL[\tilde{X} , G]	PSPACE-c	NP-c	coNP-c	coNP-c	EXPTIME-c	Π_2^P -c

Table 2: Table with complexity results on the safety fragments of LTL (the results in grey are proved in this paper).

The results stated by Theorems 4 and 5, which are summarised in Tables 1 and 2, show a surprising (a)symmetry in the complexity of the satisfiability problem along two different dimensions.

Moving from coSafety-LTL to LTL[X, F], either on infinite or on finite traces, the complexity of satisfiability changes from PSPACE-complete to NP-complete. This comes from a linear-size model property known for LTL[X, F] on infinite traces (Sistla and Clarke 1985, Lemma 3.6), which allows us to guess (nondeterministically) a candidate model and then check it in polynomial time. Instead, thanks to the *until* (U) operator in coSafety-LTL, which combines an existential quantification over time points with a bounded universal one, we are able to encode LTL formulas interpreted over finite traces with coSafety-LTL formulas that, by means of the *until* modality, can hook the final state of a finite trace and simulate the universal temporal modalities of LTL (like the *globally*) by means of the universal part of the *until*.

It is worth noticing that, being without universal temporal operators (that is, \tilde{X} , G, and R), coSafety-LTL, F(pLTL), and LTL[X, F] formulas cannot detect any difference between satisfiability on finite and on infinite traces, since any satisfying finite trace can be arbitrarily extended to an infinite model, and any satisfying infinite trace can be suitably contracted to a satisfying finite prefix. In fact, we will prove that coSafety-LTL, LTL[X, F], and F(pLTL) are suffix independent logics (Definition 5), and we will use Theorem 1 to prove that their complexities do not change when considering finite or infinite traces.

In contrast to cosafety fragments, the complexity of safety fragments significantly changes going from infinite to finite traces: while satisfiability is PSPACE-complete on infinite traces, it is NP-complete on finite traces. This is because the *weak next* (\tilde{X}) operator, available in Safety-LTL

and LTL[\tilde{X} , G], behaves on infinite traces exactly as a *strong next* (X), which, together with the *globally* (G) or the *release* (R) operators, can encode computations of Turing machines with a polynomial tape (cf. Gabbay, Hodkinson, and Reynolds 1994, Theorem 15.8.1). Instead, on finite traces, the combination of \tilde{X} and G cannot force a trace to have more than one state. In fact, for any safety fragment interpreted over finite traces, by Theorem 2, we have that any formula of these logics is satisfiable if and only if it has a model of length 1, which leads to the NP complexity.

Theorems 4 and 5 are proved in the rest of this section.

Complexity on Infinite Traces

We begin with the proof of Theorem 4, proving first the PSPACE-completeness of satisfiability on infinite traces for coSafety-LTL and F(pLTL).

We start from coSafety-LTL. To prove PSPACE-hardness, we reduce the satisfiability problem of LTL over finite traces, which is PSPACE-complete (De Giacomo and Vardi 2013), to the satisfiability of coSafety-LTL over infinite traces. For any formula $\phi \in \text{LTL}$, we will define a formula $g(\phi) \in \text{coSafety-LTL}$ such that: (i) the size of $g(\phi)$ is polynomial in the size of ϕ ; (ii) ϕ is satisfiable over finite traces if and only if $g(\phi)$ is satisfiable over infinite traces. The rationale is to introduce a fresh proposition letter e that is supposed to represent, in an infinite state sequence, the end of a finite trace. We first define a transformation $f(\cdot)$ from LTL to coSafety-LTL formulas inductively as follows:

$$\begin{aligned}
f(p) &:= p, & \text{for any } p \in \Sigma, \\
f(\neg p) &:= \neg p, & \text{for any } p \in \Sigma, \\
f(\phi_1 \wedge \phi_2) &:= f(\phi_1) \wedge f(\phi_2), \\
f(\phi_1 \vee \phi_2) &:= f(\phi_1) \vee f(\phi_2), \\
f(X\phi_1) &:= X(\neg e \wedge f(\phi_1)),
\end{aligned}$$

$$\begin{aligned}
f(\tilde{X}\phi_1) &:= X(\neg e \rightarrow f(\phi_1)), \\
f(\phi_1 \text{ U } \phi_2) &:= (\neg e \wedge f(\phi_1)) \text{ U } (\neg e \wedge f(\phi_2)), \\
f(\phi_1 \text{ R } \phi_2) &:= ((\neg e \wedge f(\phi_2)) \text{ U } (\neg e \wedge f(\phi_2) \wedge X(e))) \vee \\
&\quad ((\neg e \wedge f(\phi_2)) \text{ U } (\neg e \wedge f(\phi_1) \wedge f(\phi_2))).
\end{aligned}$$

On finite and infinite traces, $f(F\phi_1)$ and $f(G\phi_1)$ can be equivalently rewritten as $(\neg e) \text{ U } (\neg e \wedge f(\phi_1))$, and $(\neg e \wedge f(\phi_1)) \text{ U } (\neg e \wedge f(\phi_1) \wedge X(e))$, respectively. Starting from $f(\cdot)$, we define the transformation $g(\cdot) : \text{LTL} \rightarrow \text{coSafety-LTL}$ as follows: for any $\phi \in \text{LTL}$, we define $g(\phi) := \neg e \wedge f(\phi) \wedge (\neg e) \text{ U } (e)$. For any $\phi \in \text{LTL}$, $g(\phi)$ is a coSafety-LTL formula and the size of $g(\phi)$ is polynomial (more precisely, linear) in the size of ϕ . The following lemma establishes the main property for using $g(\cdot)$ as an hardness reduction.

Lemma 1. *For any $\phi \in \text{LTL}$, it holds that ϕ is satisfiable over finite traces iff $g(\phi)$ is satisfiable over infinite traces.*

Using Lemma 1, we can easily prove the following result.

Lemma 2. *The satisfiability problem for coSafety-LTL on infinite traces is PSPACE-complete.*

Proof. (Membership) Immediate from the fact that $\text{coSafety-LTL} \subseteq \text{LTL}$ and that LTL satisfiability on infinite traces is PSPACE-complete (Sistla and Clarke 1985).

(Hardness) Immediate from Lemma 1 and the fact that the size of $g(\phi)$ is polynomial in the size of ϕ . \square

We now prove the PSPACE-completeness for the satisfiability problem over infinite traces of the F(pLTL) fragment. The hardness proof is based on the simple consideration that any formula ϕ of pLTL is satisfiable (over finite traces) if and only if the formula $F(\phi)$ is satisfiable over infinite (or finite) traces. The PSPACE-hardness follows from the fact that satisfiability of pLTL is PSPACE-complete.¹

Lemma 3. *The satisfiability problem for F(pLTL) on infinite traces is PSPACE-complete.*

From Lemmas 2 and 3 and Proposition 3, it follows that the validity problem of Safety-LTL and G(pLTL) over infinite traces is PSPACE-complete.

We now focus on complexity of satisfiability for Safety-LTL, G(pLTL), and $\text{LTL}[\tilde{X}, G]$ on infinite traces, showing that all these problem are also PSPACE-complete.

PSPACE-completeness of $\text{LTL}[\tilde{X}, G]$ follows from the same proof as (Cimatti et al. 2021, Th. 5.2 Cor. 5.1) or, alternatively, by adapting the proof by (Gabbay, Hodkinson, and Reynolds 1994, Thm. 15.8.1) or the proof by (Artale et al. 2014, Thm. 4.4). PSPACE-completeness of the validity problem for F(pLTL) over infinite traces follows from Proposition 3. Moreover, since $\text{LTL}[\tilde{X}, G]$ is a syntactic fragment of Safety-LTL, it immediately follows that satisfiability (resp., validity) of Safety-LTL (resp., coSafety-LTL) over infinite traces is PSPACE-complete.

¹To see this, observe that any formula ϕ of LTL is satisfiable over finite traces iff ϕ' is satisfiable, where ϕ' is obtained from ϕ by replacing each X (resp. U) operator with Y (resp. S). The PSPACE-completeness follows from the fact that satisfiability over finite traces is PSPACE-complete (De Giacomo and Vardi 2013).

Finally, we have to prove that satisfiability of G(pLTL) is PSPACE-complete on infinite traces. To prove it, we show that the validity problem for F(pLTL) is PSPACE-complete: PSPACE-completeness of satisfiability of G(pLTL) then follows from Proposition 3. As in Lemma 3, the validity of F(pLTL) can be reduced to the validity of pLTL.

Lemma 4. *The validity problem for F(pLTL) on infinite traces is PSPACE-complete.*

From Lemma 4 and Proposition 3, it follows that G(pLTL) satisfiability on infinite traces is PSPACE-complete.

Complexity on Finite Traces

We now move to the proof of Theorem 5. We first show the PSPACE-completeness of the satisfiability (resp. validity) problem of coSafety-LTL and F(pLTL) (resp. Safety-LTL and G(pLTL)) over finite traces. To this goal, we first prove that coSafety-LTL and F(pLTL) are suffix independent. We will use this result, along with Theorem 1, to transfer the complexity of satisfiability from infinite to finite traces (cf. also Cimatti et al. (2022), Lemma 1, and Artale, Mazzullo, and Ozaki (2022), Lemma 4.11).

Lemma 5. *coSafety-LTL and F(pLTL) are suffix independent.*

From Theorem 1, we obtain the following corollary.

Corollary 1. *The satisfiability problem over finite traces of coSafety-LTL and F(pLTL) is PSPACE-complete.*

By Proposition 3, it follows that the validity problems of Safety-LTL and G(pLTL) are PSPACE-complete.

It is worth noticing that Theorem 1 does not work for safety fragments of LTL: for example, the formula $G(\tilde{X}\perp)$ is satisfiable over finite traces but unsatisfiable over infinite traces. As a matter of fact, below we show that the complexity of Safety-LTL, G(pLTL) and $\text{LTL}[\tilde{X}, G]$ satisfiability lowers down to NP-complete under finite trace semantics. Indeed, since Safety-LTL, G(pLTL) and $\text{LTL}[\tilde{X}, G]$ are safety fragments of LTL, from Theorem 2, it follows that any satisfiable formula of these fragments has a model of length 1.

Consequently, we can give a nondeterministic algorithm that, in polynomial time, solves the satisfiability of a formula $\phi \in \mathbb{L}$, with $\mathbb{L} \in \{\text{LTL}[\tilde{X}, G], \text{Safety-LTL}, \text{G(pLTL)}\}$. It simply suffices to guess an assignment for the initial state of a candidate trace and check if it satisfies ϕ . If such an assignment is found, then it means that ϕ is satisfiable, otherwise, by Theorem 2, ϕ is unsatisfiable. This proves the membership of Safety-LTL, $\text{LTL}[\tilde{X}, G]$, and G(pLTL) to NP. The hardness simply follows from a reduction of the SAT problem.

Lemma 6. *The satisfiability problem on finite traces for Safety-LTL, $\text{LTL}[\tilde{X}, G]$, and G(pLTL) is NP-complete.*

By Proposition 3, the validity problem on finite traces for coSafety-LTL, $\text{LTL}[X, F]$, and F(pLTL) is coNP-complete.

5 Complexity of Realizability

In this section, we study the complexity of the realizability problem for the (co)safety fragments of LTL that we considered in the previous section. The following theorems sum up our results on realizability over infinite and finite traces.

Theorem 6. *The realizability problem over infinite traces is*

- 2EXPTIME-complete for Safety-LTL, coSafety-LTL;
- EXPTIME-complete for $G(\text{pLTL})$, $F(\text{pLTL})$, $\text{LTL}[\tilde{X}, G]$, and $\text{LTL}[X, F]$.

Theorem 7. *The realizability problem over finite traces is*

- 2EXPTIME-complete for coSafety-LTL;
- EXPTIME-complete for $F(\text{pLTL})$, $\text{LTL}[X, F]$;
- Π_2^P -complete for Safety-LTL, $G(\text{pLTL})$, and $\text{LTL}[\tilde{X}, G]$.

Complexity on Infinite Traces

We first prove the 2EXPTIME-completeness of coSafety-LTL realizability on infinite (and finite) traces. To show hardness, we consider the realizability problem of LTL over finite traces, which is 2EXPTIME-complete (De Giacomo and Vardi 2015). For any formula ϕ of LTL, we consider the formula $g(\phi)$ as defined in the previous section, and we define the uncontrollable variable \mathcal{U}' of $g(\phi)$ as the uncontrollable variables of ϕ , and the controllable variables \mathcal{C}' of $g(\phi)$ as the set of controllable variables of ϕ and e . The following lemma establishes the equirealizability between ϕ (over finite traces) and $g(\phi)$ over infinite traces.

Lemma 7. *For any $\phi \in \text{LTL}$, it holds that ϕ is realizable over finite traces iff $g(\phi)$ is realizable over infinite traces.*

We use Lemma 7 as the core of a reduction from realizability of LTL over infinite traces to realizability of coSafety-LTL over finite traces, thus proving the following.

Lemma 8. *The realizability problem over infinite traces for coSafety-LTL is 2EXPTIME-complete.*

We now study the complexity of $\text{LTL}[\tilde{X}, G]$ and $G(\text{pLTL})$. Interestingly, for these two fragments the realizability problem over infinite traces is EXPTIME-complete. In fact, as described in (De Giacomo et al. 2021; Cimatti et al. 2021), for any formula ϕ in $\text{LTL}[\tilde{X}, G]$ or in $G(\text{pLTL})$, there exists (and can be actually built effectively) a deterministic finite automaton (DFA) $\mathcal{A}(\neg\phi)$ such that: (i) its language is exactly the set of bad prefixes of ϕ ; and (ii) its size is singly exponential in the size of ϕ .² Then, realizability can be solved on top of $\mathcal{A}(\neg\phi)$ by checking whether Controller can force the game to never visit a final state of the automaton. This kind of games, called *safety games*, can be solved in linear time. It follows that $G(\text{pLTL})$ and $\text{LTL}[\tilde{X}, G]$ realizability (over infinite traces) belongs to EXPTIME.

The EXPTIME-hardness of $\text{LTL}[\tilde{X}, G]$ follows from (Cimatti et al. 2021, Th. 5.2, Cor.5.1). The EXPTIME-hardness of $G(\text{pLTL})$ realizability over infinite words can be proved in a similar way as for the $\text{LTL}[\tilde{X}, G]$ case: for any infinite corridor tiling game \mathcal{T} (Chlebus 1986), we build a corresponding $\text{LTL}[X, G]$ formula ϕ such that \mathcal{T} admits a strategy iff ϕ is realizable. It is worth noticing that this encoding can be derived from the one of $\text{LTL}_{\text{EBR}+\text{P}}$ with

²We recall that, in the general case, the construction of a DFA starting from an LTL formula interpreted over finite traces requires two steps, each introducing an exponential blowup in the worst case: (i) the transformation of the LTL formula into a non-deterministic finite automaton (NFA); (ii) the determinization of the NFA through the classic subset construction.

no bounded operators (Cimatti et al. 2021) by using the Υ operators instead of X .

Lemma 9. *The realizability problem over infinite traces of $\text{LTL}[\tilde{X}, G]$ and $G(\text{pLTL})$ is EXPTIME-complete.*

We now prove a lemma that allows us to dualize the complexities for realizability (over infinite traces) we have found so far for coSafety-LTL, $\text{LTL}[\tilde{X}, G]$, and $G(\text{pLTL})$ to Safety-LTL, $\text{LTL}[X, F]$, and $F(\text{pLTL})$, respectively. The following lemma can be considered as the version of Proposition 3 for realizability.

Lemma 10. *Let \mathbb{L} be coSafety-LTL (resp. $\text{LTL}[X, F]$, resp. $F(\text{pLTL})$) and let \mathbb{L}' be Safety-LTL (resp. $\text{LTL}[\tilde{X}, G]$, resp. $G(\text{pLTL})$). For a complexity class C , the realizability problem over infinite traces for \mathbb{L} is C -complete iff the realizability problem over infinite traces for \mathbb{L}' is co C -complete.*

The rationale behind Lemma 10 is that realizability games are zero-sum games (Jacobs et al. 2017): Controller has a winning strategy for ϕ iff Environment has not a winning strategy for $\neg\phi$. Crucially, the existence of a winning strategy of Environment for $\neg\phi$ can be checked with classical realizability: it suffices to swap the controllable variables of $\neg\phi$ with the uncontrollable ones, and *vice versa*, and to codify in the formula the fact that Environment player has to play as the second player. Lemma 10, together with Lemmas 8 and 9, implies the following complexity results³.

Lemma 11. *The realizability problem over infinite traces for Safety-LTL (resp. $F(\text{pLTL})$ and $\text{LTL}[X, F]$) is 2EXPTIME-complete (resp. EXPTIME-complete).*

Complexity on Finite Traces

It is simple to see that Theorem 1 implies that, for any formula ϕ of coSafety-LTL, $F(\text{pLTL})$ or $\text{LTL}[X, F]$, ϕ is realizable over infinite traces iff ϕ is realizable over finite traces. Therefore, we have that the realizability problem over finite traces of coSafety-LTL, $F(\text{pLTL})$ and $\text{LTL}[X, F]$ is 2EXPTIME-complete, EXPTIME-complete, and EXPTIME-complete, respectively.

We prove that, similarly for the case of satisfiability, the complexity of safety fragments for the realizability problem significantly decreases when passing from infinite to finite traces. In particular, we prove that realizability over finite traces of $\text{LTL}[\tilde{X}, G]$, Safety-LTL and $G(\text{pLTL})$ is Π_2^P -complete. We first prove the following small model property (analogous to Theorem 2 for satisfiability), which follows from Theorem 3.

Lemma 12. *Let $\mathbb{L} \in \{\text{LTL}[\tilde{X}, G], \text{Safety-LTL}, G(\text{pLTL})\}$. Any $\phi \in \mathbb{L}$ is realizable on finite traces iff there exists a strategy $s : (2^{\mathcal{U}})^+ \rightarrow (2^{\mathcal{C}})$ such that $\text{res}(s, \mathcal{U})_0 \models \phi$, for any $\mathcal{U} \in (2^{\mathcal{U}})^\omega$.*

Lemma 12 allows for the following algorithm deciding the realizability over finite traces of $\text{LTL}[\tilde{X}, G]$, Safety-LTL and $G(\text{pLTL})$: for any ϕ in these fragments, check the existence of a strategy that satisfies ϕ in one step; if it exists, ϕ is realizable; otherwise, by Lemma 12, it is unrealizable.

³Note that this contradicts Arteche and Hermo (2021), who acknowledged a flaw in their article.

The existence of a strategy implementing ϕ in one step amounts to the check of satisfiability of a *Quantified Boolean Formula* with one quantifier alternation (2QBF), which is a Π_2^P -complete problem (Kleine Büning and Bubeck 2009). In the following we describe the algorithm.

We start with $\mathbb{L} \in \{\text{LTL}[\tilde{X}, G], \text{Safety-LTL}\}$. Let $\phi \in \mathbb{L}$ and let $\mathcal{U} = \{u_1, \dots, u_m\}$ (resp. $\mathcal{C} = \{c_1, \dots, c_n\}$) be the set of uncontrollable (resp. controllable) variables of ϕ .

1. expand the temporal operators of ϕ in the classical fashion ($G\phi$ is expanded in $\phi \wedge \tilde{X}G\phi$ and $\phi_1 R \phi_2$ is expanded in $(\phi_1 \wedge \phi_2) \vee (\phi_2 \wedge \tilde{X}(\phi_1 R \phi_2))$); the formula obtained in this way is a Boolean combination of proposition atoms or formulas of type $\tilde{X}\phi$;
2. replace each formula of type $\tilde{X}\phi$ with \top ; the resulting formula, that we call ϕ' , is a Boolean formula;
3. check the satisfiability of $\forall u_1 \dots \forall u_m \exists c_1 \dots \exists c_n . \phi'$, which is a 2QBF formula.

For $\phi \in G(\text{pLTL})$, the method is the same: for any ϕ of type $G(\alpha)$, we drop the G operator, we expand the past temporal operators in α and we replace each subformula of type $\tilde{Y}\phi$ (resp. $Y\phi$) with \top (resp. \perp). This gives us the Π_2^P -membership of $\text{LTL}[\tilde{X}, G]$, Safety-LTL and $G(\text{pLTL})$ realizability over finite traces. The Π_2^P -hardness comes directly from the Π_2^P -hardness of 2QBF.

Lemma 13. *The realizability problem over finite traces of $\text{LTL}[\tilde{X}, G]$, Safety-LTL and $G(\text{pLTL})$ is Π_2^P -complete.*

6 Discussion

The complexity gap of satisfiability for the safety fragments when moving from infinite to finite traces is worth discussing. To some extent, this shows that reducing the problem to considering prefixes of an ω -language, in the worst case, does not affect the complexity (in fact, on infinite traces, the satisfiability problem for all the fragments, except $\text{LTL}[X, F]$, has the same complexity as for full LTL). On the contrary, considering the prefixes of a language of finite words can dramatically decrease the complexity.

In the case of infinite trace semantics, in contrast to what happens for satisfiability, considering safety properties can decrease the worst-case complexity of realizability with respect to full LTL (Lemmas 9 and 11). This is due to the crucial role that *determinism* has in realizability. Indeed, realizability is (almost always) solved by playing a game over an automaton, also called *arena*, whose solution requires a deterministic representation of the arena. Therefore, reducing to reasoning over finite words (the main advantage of considering (co)safety properties) can be exploited by realizability algorithms, e.g. by building a deterministic automaton for a pLTL formula with only single exponential blowup. On the contrary, satisfiability is not able to exploit determinism to improve worst-case complexity, since it can be solved simply as the reachability of a final state in a (possibly non-deterministic) automaton corresponding to the formula. In other words, determinization is not necessary for satisfiability, and indeed coSafety-LTL and F(pLTL) share the same complexity for satisfiability.

Consider now the difference between the complexity of realizability of coSafety-LTL and F(pLTL) (or equivalently

of Safety-LTL and $G(\text{pLTL})$). Despite having the same expressive power (recall Proposition 1), the complexity is significantly lower if the formula is given in the form F(pLTL). This difference has one of these two consequences:

- either coSafety-LTL can be exponentially more succinct than F(pLTL), *i.e.*, there exists a formula $\phi \in \text{coSafety-LTL}$ such that, for all $\phi' \in F(\text{pLTL})$, if $\phi \equiv_I \phi'$ then $\phi' \in \mathcal{O}(2^{|\phi|})$;
- or there exists an algorithm of exponential running time such that, given any $\phi \in \text{coSafety-LTL}$, outputs an equivalent formula $\phi' \in F(\text{pLTL})$ with $|\phi'| \in \mathcal{O}(\text{poly}(|\phi|))$.

Clearly, exactly one of the two points can be true. We conjecture the first one to be true, but the question is still open.

As already noted in (De Giacomo et al. 2021), results on computation tree logic and alternating-time temporal logic satisfiability (Kupferman, Pnueli, and Vardi 2012; Bozzelli, Murano, and Sorrentino 2020) could be adapted to show the EXPTIME-membership of realizability for F(pLTL). It is unclear, however, how to use these results to address the EXPTIME lower bound. We also remark that our result on the 2EXPTIME-completeness of Safety-LTL realizability shows the optimality of the algorithm in (Zhu et al. 2017).

Fionda and Greco (2018) study the complexity of satisfiability for fragments of LTL over finite traces, with X, F and G as the only available temporal modalities, by imposing several syntactical restrictions and proving a linear-length model property for some of such fragments. Our study considers (together with U and R) also the role of the \tilde{X} operator, which is crucial when negation is applied only to propositional atoms. In addition, we prove that for *all* safety languages of finite words, there is a *constant-size model property*, allowing one to consider only the first state of a model.

7 Conclusions

In this paper, we studied the complexity of the (co)safety fragment of LTL for the problems of satisfiability, validity, and realizability, both over infinite and finite trace semantics. In particular, we considered three cosafety fragments (coSafety-LTL, F(pLTL), and $\text{LTL}[X, F]$) and their dual safety fragments (Safety-LTL, $G(\text{pLTL})$, and $\text{LTL}[\tilde{X}, G]$).

Our results show that: (i) for the cosafety fragment, the complexities never change when passing from infinite to finite trace semantics; (ii) on the contrary, for the safety fragment, considering finite trace semantics can significantly decrease the complexity of both satisfiability and realizability; (iii) for realizability, past operators play a crucial role; *e.g.*, by using the $G(\text{pLTL})$ fragment one can solve realizability in singly exponential time while being able to express all safety properties definable in LTL.

Model-checking is central in the field of temporal logic. A careful analysis of its complexity for the fragments that we considered in this paper is an interesting future development.

Finally, our conjecture that coSafety-LTL can be exponentially more succinct than F(pLTL) surely deserves an answer. More generally, a careful study of the succinctness of all fragments (in particular the ones that are expressively equivalent) seems a promising direction.

References

- Artale, A.; Geatti, L.; Gigante, N.; Mazzullo, A.; and Montanari, A. 2022. Complexity of Safety and coSafety Fragments of Linear Temporal Logic. *arXiv preprint*, abs/2211.14913.
- Artale, A.; Kontchakov, R.; Ryzhikov, V.; and Zakharyashev, M. 2014. A Cookbook for Temporal Conceptual Data Modelling with Description Logics. *ACM Trans. Comput. Log.*, 15(3): 25:1–25:50.
- Artale, A.; Mazzullo, A.; and Ozaki, A. 2019. Do You Need Infinite Time? In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI-19)*. AAAI Press.
- Artale, A.; Mazzullo, A.; and Ozaki, A. 2022. First-order Temporal Logic on Finite Traces: Semantic Properties, Decidable Fragments, and Applications. *arXiv preprint*, abs/2202.00610.
- Arteche, N.; and Hermo, M. 2021. On the Complexity of Realizability for Safety LTL and Related Subfragments. *arXiv preprint*, abs/2112.14102.
- Bozzelli, L.; Murano, A.; and Sorrentino, L. 2020. Alternating-time temporal logics with linear past. *Theor. Comput. Sci.*, 813: 199–217.
- Chang, E. Y.; Manna, Z.; and Pnueli, A. 1992. Characterization of Temporal Property Classes. In Kuich, W., ed., *Proceedings of the 19th International Colloquium on Automata, Languages and Programming (ICALP-92)*, 474–486. Springer.
- Chlebus, B. S. 1986. Domino-tiling games. *Journal of Computer and System Sciences*, 32(3): 374–392.
- Cimatti, A.; Geatti, L.; Gigante, N.; Montanari, A.; and Tonetta, S. 2021. Extended bounded response LTL: A new safety fragment for efficient reactive synthesis. *Formal Methods in System Design*, 1–49.
- Cimatti, A.; Geatti, L.; Gigante, N.; Montanari, A.; and Tonetta, S. 2022. A first-order logic characterisation of safety and co-safety languages. In *Proceedings of the 25th International Conference on Foundations of Software Science and Computation Structures (FOSSACS-22)*, 244–263. Springer.
- De Giacomo, G.; De Masellis, R.; and Montali, M. 2014. Reasoning on LTL on Finite Traces: Insensitivity to Infinity. In *Proceedings of the 28th National Conference on Artificial Intelligence (AAAI-14)*, 1027–1033. AAAI Press.
- De Giacomo, G.; Di Stasio, A.; Fuggitti, F.; and Rubin, S. 2021. Pure-past linear temporal and dynamic logic on finite traces. In *Proceedings of the 29th International Conference on International Joint Conferences on Artificial Intelligence (IJCAI-21)*, 4959–4965.
- De Giacomo, G.; and Vardi, M. Y. 2013. Linear Temporal Logic and Linear Dynamic Logic on Finite Traces. In Rossi, F., ed., *Proceedings of the 23rd International Joint Conference on Artificial Intelligence (IJCAI-13)*, 854–860. IJCAI/AAAI.
- De Giacomo, G.; and Vardi, M. Y. 2015. Synthesis for LTL and LDL on Finite Traces. In Yang, Q.; and Wooldridge, M. J., eds., *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI-15)*, 1558–1564. AAAI Press.
- Fionda, V.; and Greco, G. 2018. LTL on Finite and Process Traces: Complexity Results and a Practical Reasoner. *J. Artif. Intell. Res.*, 63: 557–623.
- Gabbay, D. M.; Hodkinson, I.; and Reynolds, M. A. 1994. *Temporal logic: Mathematical Foundations and Computational Aspects*, volume 1. Clarendon Press.
- Gabbay, D. M.; Kurucz, A.; Wolter, F.; and Zakharyashev, M. 2003. *Many-dimensional Modal Logics: Theory and Applications*, volume 148 of *Studies in Logic and The Foundations of Mathematics*. Elsevier.
- Jacobs, S.; Bloem, R.; Brenguier, R.; Ehlers, R.; Hell, T.; Könighofer, R.; Pérez, G. A.; Raskin, J.; Ryzhyk, L.; Sankur, O.; Seidl, M.; Tentrup, L.; and Walker, A. 2017. The first reactive synthesis competition (SYNTCOMP-14). *Int. J. Softw. Tools Technol. Transf.*, 19(3): 367–390.
- Kleine Büning, H.; and Bubeck, U. 2009. Theory of Quantified Boolean Formulas. In *Handbook of Satisfiability*, volume 185 of *Frontiers in Artificial Intelligence and Applications*, 735–760. IOS Press.
- Kupferman, O.; Pnueli, A.; and Vardi, M. Y. 2012. Once and for all. *J. Comput. Syst. Sci.*, 78(3): 981–996.
- Kupferman, O.; and Vardi, M. Y. 2001. Model checking of safety properties. *Formal Methods in System Design*, 19(3): 291–314.
- Pnueli, A. 1977. The temporal logic of programs. In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science (SFCS-77)*, 46–57. IEEE.
- Pnueli, A.; and Rosner, R. 1989. On the synthesis of an asynchronous reactive module. In *Proceedings of the International Colloquium on Automata, Languages, and Programming (ICALP-89)*, 652–671. Springer.
- Rosner, R. 1992. *Modular synthesis of reactive systems*. Ph.D. thesis, Weizmann Institute of Science.
- Safra, S. 1988. On the Complexity of omega-Automata. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science (FOCS-88)*, 319–327.
- Sistla, A. P. 1985. On characterization of safety and liveness properties in temporal logic. In *Proceedings of the Fourth Annual ACM Symposium on Principles of Distributed Computing (PODC-85)*, 39–48.
- Sistla, A. P. 1994. Safety, liveness and fairness in temporal logic. *Formal Aspects of Computing*, 6(5): 495–511.
- Sistla, A. P.; and Clarke, E. M. 1985. The complexity of propositional linear temporal logics. *Journal of the ACM (JACM)*, 32(3): 733–749.
- Thomas, W. 1988. Safety-and liveness-properties in propositional temporal logic: characterizations and decidability. *Banach Center Publications*, 1(21): 403–417.
- Vardi, M. Y.; and Wolper, P. 1986. An automata-theoretic approach to automatic program verification. In *Proceedings of the First Symposium on Logic in Computer Science*, 322–331. IEEE Computer Society.

Vardi, M. Y.; and Wolper, P. 1994. Reasoning about infinite computations. *Information and Computation*, 115(1): 1–37.

Zhu, S.; Tabajara, L. M.; Li, J.; Pu, G.; and Vardi, M. Y. 2017. A Symbolic Approach to Safety LTL Synthesis. In Strichman, O.; and Tzoref-Brill, R., eds., *Proceedings of the 13th International Haifa Verification Conference*, 147–162. Springer.