
EFFECTIVE CHARACTERIZATIONS OF SIMPLE FRAGMENTS OF TEMPORAL LOGIC USING CARTON–MICHEL AUTOMATA

SEBASTIAN PREUGSCHAT AND THOMAS WILKE

Christian-Albrechts-Universität zu Kiel, 24098 Kiel, Germany
e-mail address: {preugschat,wilke}@ti.informatik.uni-kiel.de

ABSTRACT. We present a framework for obtaining effective characterizations of simple fragments of future temporal logic (LTL) with the natural numbers as time domain. The framework is based on a form of strongly unambiguous automata, also known as prophetic automata or complete unambiguous Büchi automata and referred to as Carton–Michel automata in this paper. These automata enjoy strong structural properties, in particular, they separate the “finitary fraction” of a regular language of infinite words from its “infinitary fraction” in a natural fashion. Within our framework, we provide characterizations of several natural fragments of temporal logic, where, in some cases, no effective characterization had been known previously, and give lower and upper bounds for their computational complexity.

1. INTRODUCTION

Ever since propositional linear-time temporal logic (LTL) was introduced into computer science by Amir Pnueli in [23] it has been a major object of research. The particular line of research we are following here is motivated by the question how each individual temporal operator contributes to the expressive power of LTL. More precisely, our objective is to devise decision procedures that determine whether a given LTL property can be expressed using a given subset of the set of all temporal operators, for instance, the subset that includes “next” and “eventually”, but not “until”.

As every LTL formula interpreted in the natural numbers (the common time domain) defines a regular language of infinite words (ω -language), the aforementioned question can be viewed as part of a larger program: classifying regular ω -languages, that is, finding effective characterizations of subclasses of the class of all regular ω -languages. Over the years, many results have been established and specific tools have been developed in this program, the most fundamental result being the one that says that a regular ω -language is star-free or, equivalently, expressible in first-order logic or in LTL if, and only if, its syntactic semigroup is aperiodic [15, 27, 21].

2012 ACM CCS: [Theory of computation]: Formal languages and automata theory—Automata over infinite objects; Logic—Modal and temporal logics .

Key words and phrases: unambiguous Büchi automata, fragments of temporal logic, temporal operators, forbidden patterns in automata.

The previous result is a perfect analogue of the same result for regular languages of finite words, that is, of the classical theorems by Schützenberger [25], McNaughton and Papert [18], and Kamp [15]. In general, the situation with infinite words is more complicated than with finite words; a good example for this is given in [8], where, for instance, tools from topology and algebra are used to settle characterization problems for ω -languages.


The first characterization of a fragment of LTL over finite linear orderings was given in [5], another one followed in [10], both following a simple and straightforward approach: to determine whether a formula is equivalent to a formula in a certain fragment, one computes the minimum reverse DFA for the corresponding regular language and verifies certain structural properties of this automaton, more precisely, one checks whether certain “forbidden patterns” do not occur. The first characterization for infinite words (concerning stutter-invariant temporal properties) [20] used sequential relations on ω -words; the second (concerning the nesting depth in the until/since operator) [30] used heavy algebraic machinery and did not shed any light on the computational complexity of the decision procedures involved. In fact, the upper bound that can be derived from this work is non-elementary.

In this paper, we describe a general, conceptually simple paradigm for characterizing fragments of LTL when interpreted in the natural numbers, combining ideas from [5, 10] for finite words with the work by Carton and Michel on unambiguous Büchi automata [3, 4]. The approach works roughly as follows. To determine whether a given formula is equivalent to a formula in a given fragment, convert the formula into what is called a “prophetic automaton” in [22], check that the automaton, when viewed as an automaton on finite words, satisfies certain properties, and check that languages of finite words derived from the accepting loops (“loop languages”) satisfy certain other properties. In other words, we reduce the original problem for ω -languages to problems for languages of finite words. We show that the approach works for all reasonable fragments of future LTL and yields optimal upper bounds for the complexity of the corresponding decision procedures for all but one fragment.

Clearly, the prophetic automaton we start out with is the output of a straightforward translation; one cannot (!) expect that it provides much information about the nature of the language recognized. When we check properties of the automaton when viewed as an automaton on finite words, we first take a quotient, which makes the automaton in some sense canonical. In addition, when we derive the loop languages (representing the infinitary part of the given language) we do this with respect to that quotient, making the loop languages canonical in some sense. This approach ensures that overall we do not analyze more or less arbitrary objects derived from the given formula, but objects (languages) representing very well the nature of the property defined.

Fragments of temporal logic have been studied from different perspectives. One question that has been raised several times is what exactly is the right fragment to specify a given system. A very general answer to this has been given by Leslie Lamport in his seminal paper, [14], on the Temporal Logic of Actions. Another question that has been worked on is how the complexity of model checking depends on the particular fragment considered; results on this can already be found in the groundbreaking paper [26], by A. Prasad Sistla and Edmund Clarke. The perspective taken in this paper is different, as pointed out above.

A note on terminology. As just explained, we work with a variant (for details, see below) of the automaton model introduced by Carton and Michel in [3, 4] and named CUBA model (**C**omplete **U**nambiguous **B**üchi **A**utomata). In [22], Pin uses “prophetic automata” to refer

Figure 1: CMA which recognizes $(a + b)^* b^\omega$

to CUBA’s. In the conference version of this paper, [24], we referred to these automata as “Carton–Michel automata” (CMA) and we stick to this terminology in this paper. At the conference, STACS 2012, Thomas Colcombet gave an invited talk on determinism, non-determinism, and unambiguity with a very broad perspective and used, justified by his broad perspective, the notion “strongly unambiguous automata” (SUA) for a somewhat weaker form of unambiguity, see also the contribution to the conference proceedings, [6].

Outline. In Section 2, we provide background on the topics relevant to this paper, in particular, CMA’s, propositional linear-time temporal logic, and its translation into CMA’s. In Section 3, we present our characterizations. In Section 4 to Section 8, we give proofs of the correctness of our characterizations, and in Section 9, we explain how our characterizations can be used effectively and deal with complexity issues. We conclude with open problems.

2. BASIC NOTATION AND BACKGROUND

2.1. Reverse Deterministic Büchi Automata. A *Büchi automaton with a reverse deterministic transition function* is a tuple (A, Q, I, \cdot, F) where


- A is a finite set of symbols,
- Q is a finite set of states,
- $I \subseteq Q$ is a set of initial states,
- \cdot is a reverse transition function $A \times Q \rightarrow Q$, and
- $F \subseteq Q$ is a set of final states.

As usual, the transition function is extended to finite words by setting $\epsilon \cdot q = q$ and $au \cdot q = a \cdot (u \cdot q)$ for $q \in Q$, $a \in A$, and $u \in A^*$. For ease in notation, we write uq for $u \cdot q$ when the transition function \cdot is clear from the context.

A run of an automaton as above on an ω -word u over A is an ω -word r over Q satisfying the condition $r(i) = u(i)r(i + 1)$ for every $i < \omega$. Such a run is called *initial* if $r(0) \in I$; it is *final* if there exist infinitely many i such that $r(i) \in F$; it is *accepting* if it is initial and final. The language of ω -words recognized by such an automaton, denoted $L(\mathbf{A})$ when \mathbf{A} stands for the automaton, is the set of ω -words for which there exists an accepting run.

2.2. Carton–Michel Automata. An automaton as above is called a *Carton–Michel automaton (CMA)* if for every ω -word over A there is exactly one final run. Such an automaton is *trim*, if every state occurs in some final run.— The original definition of Carton and Michel in [3, 4] is slightly different, but for trim automata—the interesting ones—the definitions coincide.

As an example, consider the automaton depicted in Figure 1, which is a CMA for the language denoted by $(a + b)^* b^\omega$. Note that we depict $p = aq$ as



An initial state has an incoming edge \rightarrow , a final state has a double circle $\textcircled{\textcircled{\cdot}}$. Note that both components in Figure 1 belong to the automaton. The right component is needed to satisfy the condition that every ω -word has a final run in the automaton.

The fundamental result obtained by Carton and Michel is the following.

Theorem 2.1 (Carton and Michel [3, 4]). *Every regular ω -language is recognized by some CMA. More precisely, every Büchi automaton with n states can be transformed into an equivalent CMA with at most $(12n)^n$ states.*

Let \mathbf{A} be a CMA over an alphabet A and $u \in A^+$. The word u is a *loop at q* if $q = uq$ and there exist $v, w \in A^*$ satisfying $vw = u$ and $wq \in F$. The set of loops at q is denoted $S(q)$. What Carton and Michel prove about loops is:

Lemma 2.2 (Carton and Michel [3, 4]). *Let \mathbf{A} be a CMA over some alphabet A . Then, for every $u \in A^+$, there is exactly one state q , denoted u^\dagger and called anchor of u , such that u is a loop at q .*

In other words, the $S(q)$'s are pairwise disjoint and $\bigcup_{q \in Q} S(q) = A^+$.

2.3. Generalized Carton–Michel Automata. A *generalized Carton–Michel automaton (GCMA)* is defined as expected. It is the same as a CMA except that the set F of final states is replaced by a set $\mathfrak{F} \subseteq 2^Q$ of final sets, just as with ordinary generalized Büchi automata. For such an automaton, a run r is final if for every $F \in \mathfrak{F}$ there exist infinitely many i such that $r(i) \in F$.

The above definitions for CMA's can all be adapted to GCMA's in a natural fashion. For instance, a word u is a loop at some state q in a GCMA if $q = uq$ and for every $F \in \mathfrak{F}$ there exist $v, w \in A^*$ such that $u = vw$ and $wq \in F$.

It is a theorem by Carton and Michel that every GCMA can be converted into an equivalent CMA:

Theorem 2.3 (Carton and Michel [3, 4]). *Let $\mathbf{A} = (A, Q, I, \cdot, \mathfrak{F})$ be a GCMA such that $|Q| = n$ and $|\mathfrak{F}| = m$. There is an equivalent CMA $\mathbf{A}' = (A', Q', I', \cdot, \mathfrak{F}')$ such that $|Q'| \leq 2^{mn}$.*

The proof of Lemma 2.2 given in [4] carries over to GCMA's without any change. Therefore, we sometimes apply the lemma in the context of GCMA's even though it is not phrased in this context.

2.4. Temporal Logic. In the following, it is understood that temporal logic refers to propositional linear-time future temporal logic where the natural numbers are used as the domain of time. For background on temporal logic, we refer to [9] and [11]. As we are dealing with automata and formal languages, we use an approach where the atomic formulas stand for symbols of an alphabet rather than propositional variables, but note that both approaches are interchangeable.

Given an alphabet A , the set of *temporal formulas over A* , denoted TL_A , is typically inductively defined by:

- (i) for every $a \in A$, the symbol a is an element of TL_A ,
- (ii) if $\varphi \in \text{TL}_A$, so is $\neg\varphi$,
- (iii) if $\varphi, \psi \in \text{TL}_A$, so are $\varphi \vee \psi$ and $\varphi \wedge \psi$,

- (iv) if $\varphi \in \text{TL}_A$, so is $\text{X}\varphi$ (“next φ ”),
- (v) if $\varphi \in \text{TL}_A$, so are $\text{F}\varphi$ and $\text{G}\varphi$ (“eventually φ ” and “always φ ”),
- (vi) if $\varphi, \psi \in \text{TL}_A$, so are $\varphi\text{U}\psi$ and $\varphi\text{R}\psi$ (“ φ until ψ ” and “ φ releases ψ ”).

Often, the operators XF (“strictly eventually”) and XG (“strictly always”) are part of the syntax of temporal logic; we view them as abbreviations of XF and XG . For instance, $\text{XF}(a \wedge \text{XG}\neg b)$ is viewed as $\text{X}(\text{F}(a \wedge \text{X}(\text{G}\neg b)))$. (Obviously, F and G can be viewed as abbreviations of $(a \vee \neg a)\text{U}$ and $(a \wedge \neg a)\text{R}$, respectively.)

Formulas of TL_A are interpreted in ω -words over A . For every such word u , we define what it means for a formula to hold in u , denoted $u \models \varphi$, where we omit the straightforward rules for Boolean connectives:

- $u \models a$ if $u(0) = a$,
- $u \models \text{X}\varphi$ if $u[1, \ast) \models \varphi$, where, as usual, $u[1, \ast)$ denotes the word $u(1)u(2)\dots$,
- $u \models \text{F}\varphi$ if there exists $i \geq 0$ such that $u[i, \ast) \models \varphi$, similarly, $u \models \text{G}\varphi$ if $u[i, \ast) \models \varphi$ for all $i \geq 0$,
- $u \models \varphi\text{U}\psi$ if there exists $j \geq 0$ such that $u[j, \ast) \models \psi$ and $u[i, \ast) \models \varphi$ for all $i < j$, similarly, $u \models \varphi\text{R}\psi$ if there exists $j \geq 0$ such that $u[j, \ast) \models \varphi$ and $u[i, \ast) \models \psi$ for all $i \leq j$ or if $u[i, \ast) \models \psi$ for all $i \geq 0$.

Clearly, a formula of the form $\neg\text{F}\varphi$ is equivalent to $\text{G}\neg\varphi$, and a formula of the form $\neg(\varphi\text{U}\psi)$ is equivalent to $\neg\varphi\text{R}\neg\psi$, which means F and G as well as U and R are dual to each other; X is self-dual.

Given a TL_A formula φ , we write $\text{L}(\varphi)$ for the set of ω -words over A where φ holds, that is, $\text{L}(\varphi) = \{u \in A^\omega : u \models \varphi\}$. This ω -language is called the *language defined by φ* .

Given TL_A formulas φ and ψ , we say φ and ψ are *equivalent*, denoted $\varphi \equiv \psi$, if $\text{L}(\varphi) = \text{L}(\psi)$ holds.

2.5. Negation Normal Form. In the later sections of this paper, we always assume that LTL formulas can be assumed to be in negation normal form, which means (ii) from above is not used. The reason that we can do so is that \neg can easily be “pushed in”, as is explained in the following lemma.

Lemma 2.4. *Let A be some alphabet, $a \in A$, and $\varphi, \psi \in \text{TL}_A$. Then:*

$$\begin{aligned} \neg a &\equiv \bigvee_{b \in A \setminus \{a\}} b, & \neg\text{X}\varphi &\equiv \text{X}\neg\varphi, \\ \neg\text{F}\varphi &\equiv \text{G}\neg\varphi, & \neg\text{G}\varphi &\equiv \text{F}\neg\varphi, \\ \neg(\varphi\text{U}\psi) &\equiv \neg\varphi\text{R}\neg\psi, & \neg(\varphi\text{R}\psi) &\equiv \neg\varphi\text{U}\neg\psi. \end{aligned}$$

Proof hints. The proofs of the individual equivalences are straightforward. Only the proof of the second one is not generic in the sense that it fails for finite words, but for infinite words, which we only consider, no problem occurs. \square

From a complexity point of view, it is important to note that when a formula is converted to negation normal form, the size of the formula does not increase much and neither does the number of its subformulas: the increase in the length is at most the number of occurrences of alphabet symbols in the formula and the increase in the number of subformulas is at most the number of alphabet symbols. These increases do not have any influence on the upper bounds we prove in later chapters.

2.6. Fragments of Temporal Logic. An *operator set* is a subset of the set of all basic temporal operators, $\{X, F, XF, U\}$. If A is an alphabet and O an *operator set*, then $TL_A[O]$ denotes all LTL formulas that can be built from A using Boolean connectives and the operators from O . We say a language $L \subseteq A^\omega$ is *O-expressible* if there is a formula $\varphi \in TL_A[O]$ such that $L(\varphi) = L$. The *O-fragment* is the set of all LTL-formulas φ such that $L(\varphi)$ is *O-expressible*.

Observe that several operator sets determine the same fragment: $\{XF\}$ and $\{F, XF\}$; $\{U\}$ and $\{F, U\}$; $\{XF, U\}$ and $\{F, XF, U\}$; $\{X, F\}$, $\{X, XF\}$ and $\{X, F, XF\}$; $\{X, U\}$ and every superset of this.

What we are aiming at are decision procedures for each fragment except for the one determined by $\{XF, U\}$.

2.7. Ehrenfeucht–Fraïssé Games for Temporal Logic. The statements of our results (Section 3.2) do not involve Ehrenfeucht–Fraïssé games (EF games), but we use them extensively in our proofs. We make use of them in Section 5.

In the following, we recall the basics of EF games for temporal logic, see [10] for details.

A play of a temporal logic EF game is played by two players, Spoiler and Duplicator, on two ω -words over some alphabet A , say u and v . The game is played in rounds, where in every round, Spoiler moves first and Duplicator replies. The basic idea is that Spoiler is trying to reveal a difference between u and v which can be expressed in temporal logic, while Duplicator is trying to show—by somehow imitating the moves of Spoiler—that there is no such difference.

There are different types of rounds, corresponding to the temporal operators considered. We explain the ones that we need:

- ▷ *X-round.* Spoiler chooses either u or v , say v , and chops off the first letter of v , that is, he replaces v by $v[1, *]$. Duplicator does the same for u .
- ▷ *F-round.* Spoiler chooses either u or v , say v , and chops off an arbitrary finite (possibly empty) prefix, that is, he replaces v by $v[i, *]$ for some $i \geq 0$. Duplicator replaces u (the other word) by $u[j, *]$ for some $j \geq 0$.
- ▷ *XF-round.* Spoiler chooses either u or v , say v , and chops off an arbitrary non-empty finite prefix, that is, he replaces v by $v[i, *]$ for some $i > 0$. Duplicator replaces u (the other word) by $u[j, *]$ for some $j > 0$.

Before the first round, $u(0)$ and $v(0)$ are compared. If they are distinct, then this is a win (an early win) for Spoiler. After each round, the same condition is verified, and, again, if the two symbols are distinct, then this is a win for Spoiler. If, by the end of a play, Spoiler hasn't won, then this play is a win for Duplicator. For a fixed n , Duplicator wins the n -round game, if Duplicator has a strategy to win it.

When only rounds are allowed that correspond to operators in a temporal operator set $O \subseteq \{X, F, XF\}$, then we speak of an *O-game*.

The fundamental property of EF games we are going to use is the following, which was essentially proved in [10].

Theorem 2.5. *Let L be a language of ω -words over some alphabet A and $O \subseteq \{X, F, XF\}$ a temporal operator set. Then the following are equivalent:*

- (A) *L is O-expressible.*
- (B) *There is some k such that for all words $u, v \in A^\omega$ with $u \in L \leftrightarrow v \in L$, Spoiler has a strategy to win the O-game on u and v within k rounds.*

2.8. From Temporal Logic to Carton–Michel Automata. Several translations from temporal logic into Büchi and generalized Büchi automata are known, see, for instance, [31, 29, 12]. Here, we follow the ideas of these papers and “observe” that the resulting automaton is a GCMA. This is supposed to be folklore,¹ but—to the best of our knowledge—has not been made precise yet.

Let $\varphi \in \text{TL}_A$ and let $\text{sub}(\varphi)$ denote the set of its subformulas. We define a GCMA $\mathbf{A}_\varphi = (A, 2^{\text{sub}(\varphi)}, I, \cdot, \mathfrak{F})$. Our goal is to construct the automaton in such a way that in the unique final run r of this automaton on a given word u the following holds for every i and every $\psi \in \text{sub}(\varphi)$:

$$u[i, *) \models \psi \quad \text{iff} \quad \psi \in r(i) . \quad (2.1)$$

First, we set $I = \{\Phi \subseteq \text{sub}(\varphi) : \varphi \in \Phi\}$, which is motivated directly by (2.1).

Second, we define $a \cdot \Phi$ to be the smallest set $\Psi \in 2^{\text{sub}(\varphi)}$ satisfying the following conditions:

- (i) if $a \in \text{sub}(\varphi)$, then $a \in \Psi$,
- (ii) if $\neg b \in \text{sub}(\varphi)$ and $b \neq a$, then $\neg b \in \Psi$,
- (iii) if $\psi \in \Psi$ and $\chi \in \Psi$, then $\psi \wedge \chi \in \Psi$,
- (iv) if $\psi \in \Psi$ or $\chi \in \Psi$, then $\psi \vee \chi \in \Psi$,
- (v) if $\psi \in \Phi$, then $\mathbf{X}\psi \in \Psi$,
- (vi) if $\psi \in \Psi$ or $\mathbf{F}\psi \in \Phi$, then $\mathbf{F}\psi \in \Psi$,
- (vii) if $\psi \in \Psi$ and $\mathbf{G}\psi \in \Phi$, then $\mathbf{G}\psi \in \Psi$,
- (viii) if $\chi \in \Psi$ or if $\psi \in \Psi$ and $\psi \mathbf{U}\chi \in \Phi$, then $\psi \mathbf{U}\chi \in \Psi$,
- (ix) if $\chi \in \Psi$ and if $\psi \in \Psi$ or $\psi \mathbf{R}\chi \in \Phi$, then $\psi \mathbf{R}\chi \in \Psi$.

This definition reflects the “local semantics” of temporal logic, for instance, $\mathbf{F}\psi$ is true now if, and only if, ψ is true now or $\mathbf{F}\psi$ is true in the next point in time. Observe, however, that the fulfillment of $\mathbf{F}\psi$ must not be deferred forever, which means that local conditions are not enough to capture the entire semantics of temporal logic. This is taken care of by the final sets.

Third, we list the subsets of $\text{sub}(\varphi)$ which belong to \mathfrak{F} :


- for every formula $\mathbf{F}\psi \in \text{sub}(\varphi)$, the set $\{\Phi \subseteq \text{sub}(\varphi) : \psi \in \Phi \text{ or } \mathbf{F}\psi \notin \Phi\}$,
- for every formula $\mathbf{G}\psi \in \text{sub}(\varphi)$, the set $\{\Phi \subseteq \text{sub}(\varphi) : \mathbf{G}\psi \in \Phi \text{ or } \psi \notin \Phi\}$,²
- for every formula $\psi \mathbf{U}\chi$, the set $\{\Phi \subseteq \text{sub}(\varphi) : \chi \in \Phi \text{ or } \psi \mathbf{U}\chi \notin \Phi\}$,
- for every formula $\psi \mathbf{R}\chi$, the set $\{\Phi \subseteq \text{sub}(\varphi) : \psi \mathbf{R}\chi \in \Phi \text{ or } \chi \notin \Phi\}$.³

Example 2.6 ($\varphi = aRb$). After trimming, the automaton \mathbf{A}_φ for the formula $\varphi = aRb$ looks as follows.

¹Personal communication of the second author with Olivier Carton: the observation can already be found in the notes by Max Michel which he handed over to Olivier Carton in the last millennium.

²In the conference version of this paper [24] we missed this clause.

³See above.



The doubly circled states form the only final state set.

Proposition 2.7. *Let A be an alphabet and $\varphi \in TL_A$. Then \mathbf{A}_φ is a GCMA and $L(\mathbf{A}_\varphi) = L(\varphi)$.*

Proof. We first show that \mathbf{A}_φ is a GCMA. To this end, let u be an ω -word over A . We show that the word r defined by (2.1), for every i and every $\psi \in \text{sub}(\varphi)$, is a final run on u and the only one.

The ω -word r is a run on u . To see this, let $i \geq 0$ be arbitrary and observe that if we define Φ and Ψ by $\Phi = \{\psi \in \text{sub}(\varphi) : u[i+1, *) \models \psi\}$ and $\Psi = \{\psi \in \text{sub}(\varphi) : u[i, *) \models \psi\}$, then the implications (i)–(ix) not only hold, but also hold in the opposite direction. That is, $r(i) = u(i) \cdot r(i+1)$ for every i , in other words, r is a run on u .

The run r is final. Obvious from the semantics of the temporal operators.

The run r is the only possible final run. A proof of this can be carried out along the lines of the proof of Theorem 5.37 in [1], where a variant of the construction from [31] is presented and proved correct. The only differences between the setting in [1] and our setting are the atomic formulas and the set of temporal operators used. In our setting, atomic formulas correspond to letters of an alphabet; in [1], atomic formulas are propositional variables. We work with a larger set of temporal operators. In the proof in [1], it is shown that an analogue of (2.1) holds for any final run, hence the run r is the only one. \square

3. GENERAL APPROACH AND INDIVIDUAL RESULTS

This section has two purposes: it explains our general approach and presents the characterizations we have found.

3.1. The General Approach. To describe our general approach, we first need to explain what we understand by the left congruence of a GCMA.

Let \mathbf{A} be a GCMA. For every $q \in Q$, let L_q denote the set of words $u \in A^*$ such that $uq \in I$. The relation $\equiv_{\mathbf{A}}$ on Q , which we call the *left congruence of \mathbf{A}* , is defined by $q \equiv_{\mathbf{A}} q'$ when $L_q = L_{q'}$. The terminology is justified:

Remark 3.1. Let \mathbf{A} be a GCMA. Then $\equiv_{\mathbf{A}}$ is a left congruence, that is, $uq \equiv_{\mathbf{A}} uq'$ whenever $u \in A^*$ and $q, q' \in Q$ are such that $q \equiv_{\mathbf{A}} q'$.

In other words, we can define the *left quotient of \mathbf{A}* with respect to $\equiv_{\mathbf{A}}$ to be the reverse semi DFA $\mathbf{A}/\equiv_{\mathbf{A}}$ given by

$$\mathbf{A}/\equiv_{\mathbf{A}} = (A, Q'/\equiv_{\mathbf{A}}, I/\equiv_{\mathbf{A}}, \circ) \quad (3.1)$$

where

- Q' is the set of all states that occur in some final run of \mathbf{A} (active states), and
- $a \circ (q/\equiv_{\mathbf{A}}) = (a \cdot q)/\equiv_{\mathbf{A}}$ for all $a \in A$ and $q \in Q'$.

As usual, the attribute “semi” refers to the fact that this automaton has no final states nor final sets.

Next, we combine the left congruence of a GCMA with its loops. The *loop language of a state q* of a GCMA \mathbf{A} is denoted $\text{LL}(q)$ and defined by

$$\text{LL}(q) = \bigcup_{q' \equiv_{\mathbf{A}} q} S(q') \quad , \quad (3.2)$$

that is, $\text{LL}(q)$ contains all loops at q and at congruent states.

Our general approach is to characterize a fragment of LTL as follows. To check whether a given formula φ is equivalent to a formula in a given fragment, we compute the GCMA \mathbf{A}_φ and check various conditions on its left quotient and its loop languages. It turns out that this is sufficient; intuitively, the left quotient accounts for the “finitary fraction” of $L(\mathbf{A}_\varphi)$, whereas the loop languages account for its “infinitary fraction”.

3.2. Characterization of the Individual Fragments. The formal statement of our main result is as follows.

Theorem 3.2. *Let A be some alphabet, φ an LTL-formula, and O a temporal operator set as listed in Table 1. Then the following are equivalent:*

- (A) *The formula φ belongs to the O -fragment.*
- (B) *The left quotient of \mathbf{A}_φ and its loop languages satisfy the respective conditions listed in Table 1. (Information on how to read this table follows.)*

Conditions on the left quotient of \mathbf{A}_φ are phrased in terms of “forbidden patterns” (also called “forbidden configurations” in [5]). To explain this, let $\mathbf{A} = (A, Q, I, \circ)$ be any reverse semi DFA. Its *transition graph*, denoted $\text{T}(\mathbf{A})$, is the A -edge-labeled directed graph (Q, E) where $E = \{(a \circ q, a, q) : a \in A, q \in Q\}$.

Now, the conditions depicted in the second column of Table 1 are to be read as follows: the displayed graph(s) do not (!) occur as subgraphs of the transition graph of the left quotient of \mathbf{A}_φ , that is, as subgraphs of $\text{T}(\mathbf{A}_\varphi/\equiv_{\mathbf{A}_\varphi})$. Vertices filled gray must be distinct, the others may coincide (even with gray ones); dashed arrows stand for non-trivial paths.

For instance, the condition for the left quotient in the case of the $\{X\}$ -fragment requires that the following is not true for $\text{T}(\mathbf{A}_\varphi/\equiv_{\mathbf{A}_\varphi})$: there exist distinct states q and q' and a word $x \in A^+$ such that $q = x \circ q$ and $q' = x \circ q'$.

Note that for the $\{X\}$ -fragment one forbidden pattern consisting of two strongly connected components is listed, whereas for the $\{F\}$ -fragment two forbidden patterns (indicated by the horizontal line) are listed.

The conditions listed in the third column of Table 1 are conditions borrowed from formal language theory, which we explain in what follows. For a word $u \in A^*$ and $k \geq 0$, we let $\text{prf}_k(u)$, $\text{sffx}_k(u)$, and $\text{occ}_k(u)$ denote the set of prefixes, suffixes, and infixes of u of length $\leq k$, respectively. For words $u, v \in A^*$, we write $u \equiv_{k+1} v$ if $\text{prf}_k(u) = \text{prf}_k(v)$, $\text{occ}_{k+1}(u) = \text{occ}_{k+1}(v)$, and $\text{sffx}_k(u) = \text{sffx}_k(v)$. A language L is called $(k+1)$ -locally testable if $u \in L \leftrightarrow v \in L$, whenever $u \equiv_k v$, and it is called *locally testable* if it is k -locally testable for some k , see [2].

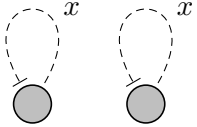
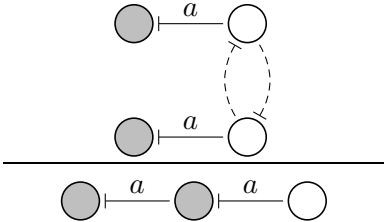
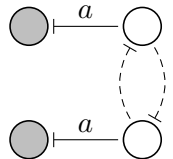
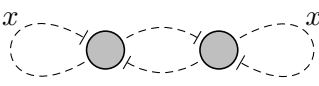

fragment	left quotient	loop languages
X		no condition
F		1-locally testable
XF		1-locally testable
X, F		locally testable
U		stutter-invariant

Table 1: Characterizations of the individual fragments of LTL

A language $L \subseteq A^+$ is *stutter-invariant* if $uav \in L \leftrightarrow uaav \in L$ holds for all $a \in A$, $u, v \in A^*$.

3.3. Proof techniques. For each fragment dealt with in Theorem 3.2, we have a separate proof, some of them are similar, others are completely different. In this section, we give a brief overview of our proofs.

For the operator set $\{X\}$, the proof is more or less a simple exercise, given that $\{X\}$ -expressibility means that there is some k such that $u \vDash \varphi$ is determined by $\text{prf}_k(u)$.

For the operator sets $\{F\}$, $\{XF\}$, and $\{X, F\}$, we use similar proofs.

For $\{U\}$, we use a theorem from [19], which says that an LTL formula over some alphabet A is equivalent to a formula in $\text{TL}_A[U]$ if the language defined by the formula is stutter-invariant, where stutter invariance is defined using an appropriate notion of stutter equivalence on ω -words.

Throughout the next sections, for ease in notation, we often write \bar{q} for $q/\equiv_{\mathbf{A}}$, where q is a state in \mathbf{A} . When $u \in A^\omega$, then $u \cdot \infty$ denotes the first state of the unique final run of \mathbf{A} on u , and $\text{inf}(u) = \{a \in A : \exists^\infty i(u(i) = a)\}$. For $a \in A$ and $u \in A^*$, $|u|_a$ denotes the number of occurrences of a in u .

4. CHARACTERIZATION OF THE $\{X\}$ -FRAGMENT

We start with the characterization of the $\{X\}$ -fragment, which is straightforward.

Theorem 4.1. *The following are equivalent for a given trim GCMA \mathbf{A} :*

- (A) $L(\mathbf{A})$ is X -expressible.
- (B) *The transition graph $T(\mathbf{A}/\equiv_{\mathbf{A}})$ does not have a subgraph of the following form (in the above sense):*



Proof. (A) implies (B): Let $L(\mathbf{A})$ be X -expressible. Let $\varphi \in \text{TL}_{\mathbf{A}}[X]$ such that $L(\mathbf{A}) = L(\varphi)$. Let $k = \text{length}(\varphi)$ where length may be any reasonable function to determine the length of a given formula φ as a natural number. Obviously for each $w \in L(\varphi)$ and $v \in A^\omega$ the following implication holds: If $\text{prf}_k(v) = \text{prf}_k(w)$ then $v \in L(\varphi)$. Let $p \in \bar{p}$ and $q \in \bar{q}$. Then there exists $u \in A^*$ with $u \cdot p \in I \Leftrightarrow u \cdot q \in I$. Let $v, v' \in A^\omega$ such that $p = v \cdot \infty$ and $q = v' \cdot \infty$. Assume that $T(\mathbf{A}/\equiv_{\mathbf{A}})$ has a subgraph of type (T1). Then $\text{prf}_k(ux^k v) = \text{prf}_k(ux^k v')$ but $ux^k v \in L(\mathbf{A}) \Leftrightarrow ux^k v' \in L(\mathbf{A})$, which is a contradiction.

We show that (B) implies (A) by contraposition. Assume $L(\mathbf{A})$ is not X -expressible. Then for every natural number k there exist $u, v \in A^\omega$ with $\text{prf}_k(u) = \text{prf}_k(v)$ and $u \in L(\mathbf{A}) \Leftrightarrow v \in L(\mathbf{A})$. Let $k \geq |Q^2|$ and u, v as described. Let r be the run of \mathbf{A} on u and s be the run of \mathbf{A} on v . Note that $r(i) \not\equiv_{\mathbf{A}} s(i)$ for every $i < k$ because $r(0) \in I$ and $s(0) \notin I$. Since $k \geq |Q^2|$ there exist $i < j < k$ with $r(i) = r(j)$ and $s(i) = s(j)$. From $\text{prf}_k(u) = \text{prf}_k(v)$ we get $u(i) \dots u(j-1) = v(i) \dots v(j-1)$ and $T(\mathbf{A}/\equiv_{\mathbf{A}})$ has a subgraph of Type (T1). \square

5. CHARACTERIZATION OF THE $\{XF\}$ -FRAGMENT

The second characterization we prove correct is the one of the $\{XF\}$ -fragment. Since every GCMA can obviously be turned into an equivalent trim GCMA, all GCMA are assumed to be trim subsequently.

We start with a refined version of Theorem 3.2 for the $\{XF\}$ -fragment.

Theorem 5.1. *The following are equivalent for a given trim GCMA \mathbf{A} :*

- (A) $L(\mathbf{A})$ is \mathcal{XF} -expressible.
- (B) (a) *The transition graph $T(\mathbf{A}/\equiv_{\mathbf{A}})$ does not have a subgraph of the following form (in the above sense):*



- (b) *For all $u, v \in A^+$ with $\text{occ}_1(u) = \text{occ}_1(v)$, it holds that $u\bar{\gamma} \equiv_{\mathbf{A}} v\bar{\gamma}$.*
- (C) (a) *The same as in (B)(a).*
 - (b) (i) *For all $u, v \in A^*$, $a \in A$, it holds that $uav\bar{\gamma} \equiv_{\mathbf{A}} uaav\bar{\gamma}$.*
 - (ii) *For all $u, v \in A^*$, $a, b \in A$, it holds that $uabv\bar{\gamma} \equiv_{\mathbf{A}} ubav\bar{\gamma}$.*

Observe that (B)(b) means that the loop languages are 1-locally testable. In other words, the above theorem implies that the characterization of the $\{\mathcal{XF}\}$ -fragment given in Theorem 3.2 is correct.

Before we get to the proof of Theorem 5.1 we provide some more notation and prove some useful lemmas.

Lemma 5.2. *Assume $\mathsf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ has a subgraph of type (T2). Then for every k there exist words $u, v \in A^\omega$ such that Duplicator wins the k -round \mathcal{XF} -game on u and v , but $u \in \mathsf{L}(\mathbf{A}) \Leftrightarrow v \in \mathsf{L}(\mathbf{A})$.*

Proof. Assume $\mathsf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ has a subgraph of type (T2). That is, there are states $\bar{p} \neq \bar{q}, \bar{r}, \bar{s}$, words $x, y \in A^+$, and a letter $a \in A$ such that $\bar{p} = a \circ \bar{r}$, $\bar{q} = a \circ \bar{s}$, $\bar{s} = y \circ \bar{r}$ and $\bar{r} = x \circ \bar{s}$. We find states r_0, r_1, \dots , and s_0, s_1, \dots such that

- $\bar{r}_i = \bar{r}$ and $\bar{s}_i = \bar{s}$ for all $i < \omega$, and
- $x \cdot s_i = r_i$ and $y \cdot r_i = s_{i+1}$ for all $i < \omega$.

Because Q is a finite set, we find $l > 0$ and i such that $r_i = r_{i+l}$. Since \mathbf{A} is trim, we find v such that $v \cdot \infty = r_i$ and u such that $ua \cdot r_i \in I$ iff $ua \cdot s_i \notin I$. This means that $ua(yx)^{lm}v \in L \Leftrightarrow uax(yx)^{lm}v \in L$ for all $m \geq 1$.

Clearly, if we choose $lm > k$, then the two resulting words cannot be distinguished in the k -round \mathcal{XF} -game. \square

Lemma 5.3. *Let \mathbf{A} be a GCMA such that $\mathsf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ does not have a subgraph of type (T2). Further, let r and s be the unique final runs of \mathbf{A} on words $u, v \in A^\omega$ and define \bar{r} and \bar{s} by $\bar{r}(i) = r(i)/\equiv_{\mathbf{A}}$ and $\bar{s}(i) = s(i)/\equiv_{\mathbf{A}}$ for all $i < \omega$.*

If $\bar{r}(0) \neq \bar{s}(0)$ and $\inf(\bar{r}) \cap \inf(\bar{s}) \neq \emptyset$, then Spoiler wins the k -round \mathcal{XF} -game on u and v where k is twice the number of states of $\mathbf{A}/\equiv_{\mathbf{A}}$.

Proof. In the following, we use SCC as an abbreviation for strongly connected component. In our context, a state which is not reachable by a non-trivial path from itself is considered to be an SCC by itself. For every $i < \omega$, let R_i and S_i be the SCC's of $\bar{r}(i)$ and $\bar{s}(i)$ in $\mathbf{A}/\equiv_{\mathbf{A}}$, respectively. Observe that because of $\inf(\bar{r}) \cap \inf(\bar{s}) \neq \emptyset$ there is some l such that the R_i 's and S_j 's are all the same for $i, j \geq l$.

Let $\mathfrak{R} = \{R_i : i > 0\}$, $\mathfrak{S} = \{S_i : i > 0\}$, $m = |\mathfrak{R}| - 1$, and $n = |\mathfrak{S}| - 1$. We show that Spoiler wins the \mathcal{XF} -game in at most $m + n$ rounds. The proof is by induction on $m + n$.

Base case. Let $m = n = 0$. Then $R_1 = S_1$. Because of the absence of (T2), we have $u(0) \neq v(0)$, and Spoiler wins instantly.

Induction step. Note that if r is the unique final run of \mathbf{A} on u , then $r[i, \ast)$ is the unique final run of \mathbf{A} on $u[i, \ast)$ for every i .

Let $m + n > 0$. If $u(0) \neq v(0)$, then Spoiler wins instantly. If $u(0) = v(0)$, we proceed by a case distinction as follows.

Case 1, $R_1 = S_1$. This is impossible because of the absence of (T2).

Case 2, $R_1 \neq S_1$, $R_1 \notin \mathfrak{S}$. Since $R_1 \notin \mathfrak{S}$ and $\inf(\bar{r}) \cap \inf(\bar{s}) \neq \emptyset$ we have $m > 0$. So there must be some $i \geq 1$ such that $\bar{r}(i) \in R_1$ and $\bar{r}(i+1) \notin R_1$. Spoiler chooses the word u and replaces u by $u[i, \ast)$.

Now Duplicator has to replace v by $v[j, \ast)$ for some $j > 0$. Since $R_1 \notin \mathfrak{S}$ we have $\bar{r}(i) \neq \bar{s}(j)$ and the induction hypothesis applies.

Case 3, $R_1 \neq S_1$, $S_1 \notin \mathfrak{R}$. Symmetric to Case 2.

Case 4, $R_1 \neq S_1$, $R_1 \in \mathfrak{S}$, and $S_1 \in \mathfrak{R}$. Impossible, because R_1 would be reachable from S_1 and vice versa, which would mean R_1 and S_1 coincide. \square

Lemma 5.4. *Let \mathbf{A} be a GCMA. Then the following are equivalent:*

- (A) *For all $u, v \in A^+$ with $\text{occ}_1(u) = \text{occ}_1(v)$, it holds that $u \vDash_{\mathbf{A}} v \vDash$.*
- (B) (a) *For all $u, v \in A^*$, $a \in A$, it holds that $uav \vDash_{\mathbf{A}} uaav \vDash$.*
 (b) *For all $u, v \in A^*$, $a, b \in A$, it holds that $uabv \vDash_{\mathbf{A}} ubav \vDash$.*

Proof. That (A) implies (B) is obvious. For the converse, let $u, v \in A^+$ with $\text{occ}_1(u) = \text{occ}_1(v)$. Let $\text{occ}_1(u) = \{a_0, a_1, \dots, a_n\}$. Now, we have

$$u \vDash_{\mathbf{A}} a_0^{|u|a_0} a_1^{|u|a_1} \dots a_n^{|u|a_n} \vDash_{\mathbf{A}} a_0^{|v|a_0} a_1^{|v|a_1} \dots a_n^{|v|a_n} \vDash_{\mathbf{A}} v \vDash ,$$

where the first and the last equivalence are obtained by iterated application of (b), and the second equivalence is obtained by iterated application of (a). \square

In what follows, we need more notation and terminology. A word $u \in A^\omega$ is an *infinite loop* at q if $q = u \cdot \infty$ and $q \in \text{inf}(r)$ where r is the unique final run of \mathbf{A} on u .

Proof of Theorem 5.1. The implication from (A) to (B)(a) is Lemma 5.2. We prove that (A) implies (B)(b) by contraposition. Assume (B)(b) does not hold, that is, there are $u, v \in A^+$ with $\text{occ}_1(u) = \text{occ}_1(v)$, and $u \vDash \not\equiv_{\mathbf{A}} v \vDash$. Then there exists $x \in A^*$ such that $x \cdot u \vDash \in I \Leftrightarrow x \cdot v \vDash \in I$, that is, $xu^\omega \in L \Leftrightarrow xv^\omega \in L$. It is easy to see that Duplicator wins the \mathbf{XF} -game on xu^ω and xv^ω for any number of rounds, which, in turn, implies L is not \mathbf{XF} -expressible.

For the implication from (B) to (A), let n be the number of states of $\mathbf{A}/\equiv_{\mathbf{A}}$. We show that whenever $u, v \in A^\omega$ such that $u \in L \Leftrightarrow v \in L$, then Spoiler wins the $2n$ -round \mathbf{XF} -game on u and v .

Assume $u, v \in A^\omega$ are such that $u \in L \Leftrightarrow v \in L$ and let r and s be the unique final runs of \mathbf{A} on u and v , respectively, and \bar{r} and \bar{s} defined as in Lemma 5.3. We distinguish two cases.

First case, $\text{inf}(u) \neq \text{inf}(v)$. Then Spoiler wins within 2 rounds.

Second case, $\text{inf}(u) = \text{inf}(v)$. Then there are i, i' and j, j' such that

- $\text{occ}_1(u[i, j]) = \text{occ}_1(v[i', j'])$,
- $u[i, *] \cdot \infty$ is an infinite loop at $u[i, j] \vDash$, and
- $v[i', *] \cdot \infty$ is an infinite loop at $v[i', j'] \vDash$.

From (B)(b), we conclude $u[i, j] \vDash \equiv_{\mathbf{A}} v[i', j'] \vDash$. As a consequence, $\text{inf}(\bar{r}) \cap \text{inf}(\bar{s}) \neq \emptyset$. Since $\bar{r}(0) \neq \bar{s}(0)$, Lemma 5.3 applies: L is \mathbf{XF} -expressible.


The equivalence between (B) and (C) follows directly from Lemma 5.4. \square

6. CHARACTERIZATION OF THE $\{\mathbf{F}\}$ -FRAGMENT

The characterization of the $\{\mathbf{F}\}$ -fragment is similar to the one of the $\{\mathbf{XF}\}$ -fragment, but a little more complicated.

Theorem 6.1. *The following are equivalent for a given trim GCMA \mathbf{A} :*

- (A) *$L(\mathbf{A})$ is \mathbf{F} -expressible.*
- (B) (a) *The transition graph $\mathsf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ does not have a subgraph of the following form (in the above sense):*



- (b) For all $u, v \in A^+$ with $\text{occ}_1(u) = \text{occ}_1(v)$ it holds that $u\bar{\jmath} \equiv_{\mathbf{A}} v\bar{\jmath}$.
- (C) (a) The same as in (B)(a).
 (b) (i) For all $u, v \in A^*$, $a \in A$ it holds that $uav\bar{\jmath} \equiv_{\mathbf{A}} uaav\bar{\jmath}$.
 (ii) For all $u, v \in A^*$, $a, b \in A$ it holds that $uabv\bar{\jmath} \equiv_{\mathbf{A}} ubav\bar{\jmath}$.

As seen above (B)(b) means that the loop languages are 1-locally testable. In other words, the above theorem implies that the characterization of the $\{\mathbf{F}\}$ -fragment given in Theorem 3.2 is correct.

Before we turn to the proof we will state some useful lemmas:

Lemma 6.2. *Assume $\mathbf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ has a subgraph of type (T2) or (T3). Then for every k there exist words $u, v \in A^\omega$ such that Duplicator wins the k -round \mathbf{F} -game on u and v , but $u \in \mathbf{L}(\mathbf{A}) \Leftrightarrow v \in \mathbf{L}(\mathbf{A})$.*

Proof. First, assume $\mathbf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ has a subgraph of type (T3). That is, there are states $\bar{p}, \bar{q}, \bar{r} \in Q/\equiv_{\mathbf{A}}$ and a symbol a such that $a \circ \bar{r} = \bar{q}$, $a \circ \bar{q} = \bar{p}$, and $\bar{p} \neq \bar{q}$. Let $r \in \bar{r}$ and define p and q by $q = a \cdot r$ and $p = a \cdot q$. Then $p \in \bar{p}$ and $q \in \bar{q}$, because $\equiv_{\mathbf{A}}$ is a left congruence.

There is some $v \in A^\omega$ such that $v \cdot \infty = r$. Further, since $\bar{p} \neq \bar{q}$, there is some $u \in A^*$ such that $u \cdot p \in I \Leftrightarrow u \cdot q \in I$. In other words, $uav \in L \Leftrightarrow uaav \in L$. Clearly, the two words cannot be distinguished in the \mathbf{F} -game.

Second, assume $\mathbf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ has a subgraph of type (T2). That is, there are states $\bar{p} \neq \bar{q}, \bar{r}, \bar{s}$ words $x, y \in A^+$ and $a \in A$ such that $\bar{p} = a \circ \bar{r}$, $\bar{q} = a \circ \bar{s}$, $\bar{s} = y \circ \bar{r}$ and $\bar{r} = x \circ \bar{s}$. We find states r_0, r_1, \dots and s_0, s_1, \dots such that

- (1) $\bar{r}_i = \bar{r}$ and $\bar{s}_i = \bar{s}$ for all i ,
- (2) $x \cdot s_i = r_i$ and $y \cdot r_i = s_{i+1}$ for all i .

Because Q is a finite set, we find $l > 0$ and i such that $r_i = r_{i+l}$. In addition, we find v such that $v \cdot \infty = r_i$ and u such that $ua \cdot r_i \in I \Leftrightarrow ua \cdot s_i \in I$. This means that $ua(yx)^{lm}v \in L \Leftrightarrow uax(yx)^{lm}v \in L$ for all $m \geq 1$.

Clearly, if we choose $lm \geq k$, then the two resulting words cannot be distinguished in the k -round \mathbf{F} -game. \square

Lemma 6.3. *Let \mathbf{A} be a GCMA such that $\mathbf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ does not have a subgraph of type (T2) or (T3). Further, let r and s be the unique final runs of \mathbf{A} on words $u, v \in A^\omega$ and define \bar{r} and \bar{s} by $\bar{r}(i) = r(i)/\equiv_{\mathbf{A}}$ and $\bar{s}(i) = s(i)/\equiv_{\mathbf{A}}$ for all $i < \omega$.*

If $\bar{r}(0) \neq \bar{s}(0)$ and $\inf(\bar{r}) \cap \inf(\bar{s}) \neq \emptyset$, then Spoiler wins the k -round \mathbf{F} -game on u and v where k is twice the number of states of $\mathbf{A}/\equiv_{\mathbf{A}}$.

Proof. Let R_i and S_j be the SCC's of $\bar{r}(i)$ and $\bar{s}(j)$ in $\mathbf{A}/\equiv_{\mathbf{A}}$, respectively.

There are i and j such that the SCC's of $\bar{r}(i')$ and $\bar{s}(j')$ for $i' \geq i$ and $j' \geq j$ are all the same.

Let $\mathfrak{R} = \{R_i : i > 0\}$, $\mathfrak{S} = \{S_i : i > 0\}$, $m = |\mathfrak{R}|$, and $n = |\mathfrak{S}|$. We show that Spoiler wins the game in at most $m + n$ rounds. The proof is by induction on $m + n$. If $u(0) \neq v(0)$ Spoiler wins instantly. Otherwise, we distinguish several cases.

Case 1, $R_1 = S_1$. This is impossible because of the absence of (T2).

Case 2, $R_1 \neq S_1$, $R_1 \notin \mathfrak{S}$. Since $R_1 \notin \mathfrak{S}$ and $\inf(\bar{r}) \cap \inf(\bar{s}) \neq \emptyset$ we have $m > 1$. So there must be some $i \geq 1$ such that $\bar{r}(i) \in R_1$ and $\bar{r}(i+1) \notin R_1$. Spoiler chooses the word u and replaces u by $u[i, *]$.

If $u(i) \neq v(0)$ Duplicator has to replace v by $v[j, *]$ for some $j > 0$ if she does not want to lose right away. The induction hypothesis applies since $\bar{r}(i) \in R_1 \notin \mathfrak{S}$ and so $\bar{r}(i) \neq \bar{s}(j)$.

If $u(i) = v(0)$ ($= u(0)$), we have to show that $\bar{r}(i) \neq \bar{s}(0)$ to be able to apply the induction hypothesis. Assume that $\bar{r}(i) = \bar{s}(0)$. Since $\bar{r}(i) \in R_1 \notin \mathfrak{S}$ and $\bar{s}(1) \in S_1$, we have $\bar{s}(0) \neq \bar{s}(1)$ and $\bar{s}(0) = v(0) \circ \bar{s}(1)$, i. e. $\bar{r}(i) = v(0) \circ \bar{s}(1)$ and $\bar{r}(i) \neq \bar{s}(1)$. The absence of (T3) leads to $\bar{r}(i) = v(0) \circ \bar{r}(i)$ and the absence of (T2) leads to $\bar{r}(0) = v(0) \circ \bar{r}(i)$. We get $\bar{r}(0) = \bar{r}(i) = \bar{s}(0)$ —a contradiction.

Case 3, $R_1 \neq S_1$, $S_1 \notin \mathfrak{R}$. Symmetric to Case 2.

Case 4, $R_1 \neq S_1$, $R_1 \in \mathfrak{S}$, and $S_1 \in \mathfrak{R}$. Impossible, because R_1 would be reachable from S_1 and vice versa, which would mean R_1 and S_1 coincide. \square

Proof of Theorem 6.1. That (A) implies (B)(a) follows from Lemma 6.2 by contraposition.

We prove that (A) implies (B)(b) by contraposition. Assume (B)(b) does not hold. Then there are $u, v \in A^+$ with $\text{occ}_1(u) = \text{occ}_1(v)$ and $u \not\equiv_{\mathbf{A}} v$. Then there exists $x \in A^*$ such that $x \cdot u \in I \Leftrightarrow x \cdot v \in I$, that is, $xu^\omega \in L \Leftrightarrow xv^\omega \in L$. Now it is easy to see that Duplicator wins the F-game on xu^ω and xv^ω for any number of rounds, which, in turn, implies L is not F-expressible.

For the implication from (B) to (A), let n be the number of states of $\mathbf{A}/\equiv_{\mathbf{A}}$. We show that whenever $u, v \in A^\omega$ such that $u \in L \Leftrightarrow v \in L$ Spoiler wins the $2n$ -round F-game on u and v .

Assume $u, v \in A^\omega$ are such that $u \in L$ and $v \notin L$. We distinguish two cases.

First case, $\inf(u) \neq \inf(v)$. Then Spoiler wins within at most 2 rounds.

Second case, $\inf(u) = \inf(v)$. Then there are i, i' and j, j' such that

- $\text{occ}_1(u[i, j]) = \text{occ}_1(v[i', j'])$,
- $u[i, *] \cdot \infty$ is an infinite loop at $u[i, j]$, and
- $v[i', *] \cdot \infty$ is an infinite loop at $v[i', j']$.

From (B)(b), we conclude $u[i, j] \equiv_{\mathbf{A}} v[i', j']$. As a consequence, Lemma 6.3 applies: L is F-expressible.

The equivalence between (B) and (C) follows directly from Lemma 5.4. \square

7. CHARACTERIZATION OF THE $\{\mathbf{X}, \mathbf{F}\}$ -FRAGMENT

The correctness proof for the characterization of $\{\mathbf{X}, \mathbf{F}\}$ -fragment follows the one for the $\{\mathbf{F}\}$ -fragment. We begin with a theorem corresponding to Theorems 6.1 and 5.1.

Theorem 7.1. *The following are equivalent for a given trim GCMA \mathbf{A} :*

- (A) $L(\mathbf{A})$ is XF-expressible.
 (B) (a) *The transition graph $T(\mathbf{A}/\equiv_{\mathbf{A}})$ does not have a subgraph of the following form (in the above sense):*



- (b) For some natural k and all $u, v \in A^+$ with $u \equiv_{k+1} v$ we have $u^\natural \equiv_{\mathbf{A}} v^\natural$, i. e., for every $q \in Q$, the set $\cup_{q' \equiv_{\mathbf{A}} q} S(q')$ is locally testable.

Before we turn to the proof we will again state some useful lemmas:

Lemma 7.2. *Assume the transition graph $\mathsf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ has a subgraph of type (T4). Then for every k there exist words $u, v \in A^\omega$ such that Duplicator wins the k -round XF-game on u and v , but $u \in L \Leftrightarrow v \in L$.*

Proof. Assume (T4) occurs in $\mathsf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$. First observe that for every state $p_i, q_j \in Q$ with $\bar{p}_i = \bar{p}$ and $\bar{q}_j = \bar{q}$ and every l it holds that

$$z^l \circ \bar{p}_i = \bar{p}, z^l \circ \bar{q}_j = \bar{q}, \bar{p} = x \circ \bar{q}_j, \bar{q} = y \circ \bar{p}_i.$$

Then, observe that for every state r there exist k and $l > 0$ such that $z^k \cdot r = z^{k+l} \cdot r$. Since k can be replaced by any larger number and l by any multiple of l , we can assume k and l are the same for all states. Let l be fixed with that properties.

Let $p_i \in Q$ with $\bar{p}_i = \bar{p}$. Since Q is finite, there exist j, m with $(xz^{2l}yz^{2l})^j \cdot p_i = (xz^{2l}yz^{2l})^{j+m} \cdot p_i$. It follows easily, that there exist $x', y', z' \in A^+$ and $p', q' \in Q$ with $p' \neq q'$, $\bar{p}' = \bar{p}$, $\bar{q}' = \bar{q}$ and

$$p' = z' \cdot p', \quad q' = z' \cdot q', \quad p' = x' \cdot q', \quad q' = y' \cdot p',$$

meaning that $\mathsf{T}(\mathbf{A})$ also has a subgraph of type (T4).

In addition, we find $u \in A^*$ such that $u \cdot p' \in I \Leftrightarrow u \cdot q' \in I$ and $v \in A^\omega$ such that $p' = v \cdot \infty$ and $q' = y' \cdot v \cdot \infty$. This means that $u((z')^n x' (z')^n y')^n (z')^n v \in L \Leftrightarrow u y' ((z')^n x' (z')^n y')^n (z')^n v \in L$ for all $n \geq 1$.

Clearly, if we choose $n > k$, then the two resulting words cannot be distinguished in the k -round XF-game. \square

Lemma 7.3. *Let \mathbf{A} be a GCMA such that $\mathsf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ does not have a subgraph of type (T4). Further, let r and s be the unique final runs of \mathbf{A} on words $u, v \in A^\omega$ and define \bar{r} and \bar{s} by $\bar{r}(i) = r(i)/\equiv_{\mathbf{A}}$ and $\bar{s}(i) = s(i)/\equiv_{\mathbf{A}}$ for all $i < \omega$.*

Assume $\bar{r}(0) \neq \bar{s}(0)$ and $\inf(\bar{r}) \cap \inf(\bar{s}) \neq \emptyset$. Let

$$\Omega = \{Q_i \subseteq Q/\equiv_{\mathbf{A}} : Q_i \text{ is an SCC of } Q/\equiv_{\mathbf{A}}\}$$

and

$$K = 2 \sum_{Q_i \in \Omega} |Q_i|^2 + 2.$$

Then Spoiler wins the K -round XF-game on u and v .

Proof. Let R_i and S_j be the SCC's of $\bar{r}(i)$ and $\bar{s}(j)$ in $\mathbf{A}/\equiv_{\mathbf{A}}$, respectively.

There are i and j such that the SCC's of $\bar{r}(i')$ and $\bar{s}(j')$ for $i' \geq i$ and $j' \geq j$ are all the same.

Let $\mathfrak{R} = \{R_i : i > 0\}$, $\mathfrak{S} = \{S_i : i > 0\}$, $m = |\mathfrak{R}|$, and $n = |\mathfrak{S}|$. We show that Spoiler wins the game in at most K rounds. The proof is by induction on the induction parameter

$$\sum_{R \in \mathfrak{R}} |R|^2 + \sum_{S \in \mathfrak{S}} |S|^2 + [\bar{r}(0) \notin R_1] + [\bar{s}(0) \notin S_1].$$

Here, $[\bar{r}(0) \notin R_1]$ yields 1 if the condition is true and 0 otherwise, similarly for $[\bar{s}(0) \notin S_1]$. Adding $[\bar{r}(0) \notin R_1]$ and $[\bar{s}(0) \notin S_1]$ makes sure that if $\bar{r}(0) \notin R_1$ or $\bar{s}(0) \notin S_1$, then a X-move decreases the induction parameter. If $u(0) \neq v(0)$ Spoiler wins instantly. Otherwise, we distinguish several cases.

Case 1, $R_1 = S_1$. Let $c = |R_1|^2$. Spoiler plays c X-rounds. If Spoiler does not win in these rounds, then $\bar{r}(c+1) \notin R_1$ or $\bar{s}(c+1) \notin S_1$ because $\mathsf{T}(\mathbf{A}/\equiv_{\mathbf{A}})$ does not have a subgraph of type (T4), and, since \mathbf{A} is co-deterministic, $\bar{r}(c) \neq \bar{s}(c)$. The induction hypothesis applies.

Case 2, $R_1 \neq S_1$, $R_1 \notin \mathfrak{S}$. Then $m > 1$ and there must be some $i \geq 1$ such that $r(i) \in R_1$ and $r(i+1) \notin R_1$. We distinguish two subcases.

Subcase 2.a, $\bar{r}(i) = \bar{s}(0)$. Spoiler plays a X-round, which means Spoiler wins right away or the game proceeds with words such that their runs start in $\bar{r}(1)$ and $\bar{s}(1)$, respectively. The induction hypothesis applies, as $s(0) \notin S_1$, see above.

Subcase 2.b, $\bar{r}(i) \neq s(0)$. Spoiler plays an F-round, chooses the word u , and replaces u by $u[i, *]$. The induction parameter decreases by this, because $|R_1| \geq 2$ or $\bar{r}(0) \notin R_1$. If Duplicator chooses to not change v , then the resulting runs start with $\bar{r}(i)$ and $\bar{s}(0)$, which are distinct. If not, then the runs start with $\bar{r}(i)$ and $\bar{s}(j)$ for some $j \geq 1$, which are states that do not belong to the same SCC and, hence, are distinct.

Case 3, $R_1 \neq S_1$, $S_1 \notin \mathfrak{R}$, and $n > 1$. Symmetric to Case 2.

Case 4, $R_1 \neq S_1$, $R_1 \in \mathfrak{S}$, and $S_1 \in \mathfrak{R}$. Impossible, because R_1 would be reachable from S_1 and vice versa, which would mean R_1 and S_1 coincide. \square

For ω -words u and v and a natural number k , we write $u \approx_{k+1} v$ if $\text{prf}_k(u) = \text{prf}_k(v)$ and $\text{occ}_{k+1}(u) = \text{occ}_{k+1}(v) = \text{inf}_{k+1}(u) = \text{inf}_{k+1}(v)$.

Remark 7.4.

- (1) \approx_k is an equivalence relation.
- (2) If $u \approx_{k+1} v$, then u and v cannot be distinguished by the k -round FX-game.

We can finally turn to the correctness proof of our characterization.

Proof of Theorem 7.1. That (A) implies (B)(a) follows from Lemma 7.2 by contraposition.

We prove that (A) implies (B)(b) by contraposition. Assume (B)(b) does not hold. Let k be a natural number. There are $u, v \in A^+$ with $u \equiv_{k+1} v$ and $u \not\equiv_{\mathbf{A}} v$. Then there exists $x \in A^*$ such that $x \cdot u \in I \leftrightarrow x \cdot v \in I$, that is, $xu^\omega \in L \leftrightarrow xv^\omega \in L$. Remark 7.4 implies Duplicator wins the k -round XF-game on u^ω and v^ω because of $u^\omega \approx_{k+1} v^\omega$. But this implies Duplicator wins the XF-game on xu^ω and xv^ω , which, in turn, implies L is not XF-expressible.

For the implication from (B) to (A), let K be as in Lemma 7.3.

We show that whenever $u, v \in A^\omega$ such that $u \in L \leftrightarrow v \in L$ Duplicator wins the $\max\{K, 2+k\}$ -round XF-game on u and v .

Assume $u, v \in A^\omega$ are such that $u \in L$ and $v \notin L$. We distinguish two cases.

First case, $\text{inf}_k(u) \neq \text{inf}_k(v)$. Then Spoiler wins within at most $2+k$ rounds.

Second case, $\text{inf}_k(u) = \text{inf}_k(v)$. Then there are i, i' and j, j' such that

- $u[i, j] \equiv_k v[i', j']$,
- $u[i, *] \cdot \infty$ is an infinite loop at $u[i, j]$, and
- $v[i', *] \cdot \infty$ is an infinite loop at $v[i', j']$.

From (B)(b), we conclude $u[i, j] \equiv_{\mathbf{A}} v[i', j']$. As a consequence, Lemma 7.3 applies: L is XF-expressible. \square

8. CHARACTERIZATION OF THE $\{U\}$ -FRAGMENT

As mentioned above, the proof for the characterization of the $\{U\}$ -Fragment uses a different approach. We begin by stating the result as a theorem.

Theorem 8.1. *The following are equivalent for a given trim GCMA \mathbf{A} :*

- (A) $L(\mathbf{A})$ is U -expressible.
- (B) (a) $L(\mathbf{A})$ is TL_A -definable.
- (b) *The transition graph $T(\mathbf{A}/\equiv_{\mathbf{A}})$ does not have a subgraph of the following form (in the above sense):*

$$\textcircled{p} \xrightarrow{a} \textcircled{q} \xrightarrow{a} \textcircled{\bar{r}} \quad (\text{T3})$$

- (c) *For all $u, v \in A^*$, $a \in A : uav \equiv_{\mathbf{A}} uaav$, i. e., for every $q \in Q$, the set $\bigcup_{q' \equiv_{\mathbf{A}} q} S(q')$ is stutter-invariant.*

The definition of *stutter-invariance* for languages of ω -words is a little different to the one for finite words. We use the definition from [19]. Two ω -words u and v over an alphabet A are called *stutter-equivalent* iff there are two infinite sequences $0 = i_0 < i_1 < i_2 < \dots$ and $0 = j_0 < j_1 < j_2 < \dots$ such that for every $k \geq 0$ $u(i_k) = u(i_k + 1) = \dots = u(i_{k+1} - 1) = v(j_k) = v(j_k + 1) = \dots = v(j_{k+1} - 1)$. With the notion of stutter-equivalence we define stutter-invariance for ω -languages. An ω -Language L over an alphabet A is said to be stutter-invariant iff for each pair u, v of stutter-equivalent words we have $u \in L \leftrightarrow v \in L$.

For the proof of the above theorem, we need a theorem from [19] which reads as follows.

Theorem 8.2. *A TL_A -definable ω -language $L \subseteq A^\omega$ is U -expressible if and only if L is stutter-invariant.*

Proof of Theorem 8.1. (A) implies (B)(b): Let $L(\mathbf{A})$ be U -expressible. By Theorem 8.2 $L(\mathbf{A})$ is stutter-invariant. Assume $T(\mathbf{A}/\equiv_{\mathbf{A}})$ has a subgraph of type (T3). Then there exist $u \in A^*$, $p \in \bar{p}$ and $q \in \bar{q}$ with $u \cdot p \in I \leftrightarrow u \cdot q \in I$. Since \mathbf{A} is trim, there exists $v \in A^\omega$ with $v \cdot \infty = q$. So we have $uav \in L(\mathbf{A}) \leftrightarrow uaav \in L(\mathbf{A})$ which means $L(\mathbf{A})$ is not stutter-invariant — a contradiction.

(A) implies (B)(c) by contraposition: Assume there are $u, v \in A^*$, $a \in A$ with $uav \not\equiv_{\mathbf{A}} uaav$. then there exists $w \in A^*$ with $w \cdot uav \in I \leftrightarrow w \cdot uaav \in I$. Hence $w(uav)^\omega \in L(\mathbf{A}) \leftrightarrow w(uaav)^\omega \in L(\mathbf{A})$ and so $L(\mathbf{A})$ is not stutter-invariant.

To prove the implication from (B) to (A) we have to show that $L(\mathbf{A})$ is TL_A -definable and stutter-invariant. Then we can apply Theorem 8.2 and the proof is complete.

First we show, that $L(\mathbf{A})$ is stutter-invariant. Let $w \in L(\mathbf{A})$ with the unique final run r and \bar{r} the factorization of r as seen above. Let $i \in \mathbf{N}$ with $r(i) \in \text{inf}(r)$ and $u = w[0, i)$ and $v = w[i, *)$. Then $r(i) = v \cdot \infty$. Since the loop languages are stutter-invariant and $r(i) \in \text{inf}(r)$, for every $v' \in A^*$, $a \in A$ and $v'' \in A^\omega$ with $v'av'' = v$ there exists $r'(i) \in Q$ with $r'(i) \equiv_{\mathbf{A}} r(i)$ and $r'(i) = v'aav'' \cdot \infty$ which means $uv'aav'' \in L(\mathbf{A})$. Since $r(i) \in \text{inf}(r)$ this argument can be applied infinitely often at once. The absence of (T3) means that for every $x \in A^*$ and every $a \in A$ the equivalence $ax \cdot r(i) \equiv_{\mathbf{A}} aax \cdot r(i)$ holds. If $u \neq \varepsilon$ let $u' \in A^*$, $a \in A$ and $u'' \in A^*$ with $u'au'' = u$. Then $u'aau'' \cdot r(i) \in I$ and $u'aau''v \in L(\mathbf{A})$. So $L(\mathbf{A})$ is stutter-invariant. \square

9. EFFECTIVENESS AND COMPUTATIONAL COMPLEXITY

To conclude, we explain how Theorem 3.2 can be used effectively. In general, we have:

Theorem 9.1. *Each of the fragments listed in Table 1 is decidable.*

Observe that for the fragment with operator set $\{F, U\}$, this is a result from [20], and for the fragment with operator set $\{X, F\}$, this is a result from [30].

Proof of Theorem 9.1. First, observe that \mathbf{A}_φ can be constructed effectively. Also, it is easy to derive the left quotient of \mathbf{A}_φ from \mathbf{A}_φ itself and DFA's for the loop languages, even minimum-state DFA's for them, simply by using any of the available minimization procedures, for instance, the one described in [13].

Second, observe that the presence of the listed forbidden patterns can be checked effectively. The reason is as follows. The test for the existence of a path between two states can be restricted to paths of length at most the number of states. The test for the existence of two loops with the same label but distinct starting states (see forbidden patterns for $\{X\}$ and $\{X, F\}$) in some semi automaton $\mathbf{A} = (A, Q, \delta)$ amounts to searching for a loop in the semi automaton $(A, Q \times Q \setminus \{(q, q) \mid q \in Q\}, \delta')$ with transition function defined by $\delta'((q, q'), a) = (\delta(q, a), \delta(q', a))$. In other words, this amounts to a search in the original automaton restricted to paths of length at most the number of states squared.

Third, the conditions on the loop languages can be checked effectively. For 1-local testability, this is because a language $L \subseteq A^*$ is not 1-locally testable if, and only if, one of the following conditions holds:

- (1) There are words $u, v \in A^*$ and there is a letter $a \in A$ such that $uav \in L \leftrightarrow uaav \in L$.
- (2) There are words $u \in A^*, v \in A^*$ and letters $a, b \in A$ such that $uabv \in L \leftrightarrow ubav \in L$.

Again, u and v can be bounded in length by the number of states. For local testability, we refer to [16], where it was shown this can be decided in polynomial time. For stutter invariance, remember that a language $L \subseteq A^*$ is not stutter-invariant if, and only if, the first from the above conditions holds. So this can be checked effectively, too. (One could also use the forbidden pattern listed.) \square

As to the computational complexity of the problems considered, we first note:

Proposition 9.2. *Each of the fragments listed in Table 1 is PSPACE-hard.*

Proof. The proof is an adaptation of a proof for a slightly weaker result given in [20].

First, recall that LTL satisfiability is PSPACE-hard for some fixed alphabet [26], hence LTL unsatisfiability for this alphabet is PSPACE-hard, too. Let A denote such an alphabet in the following.

Second, let $c, d,$ and e be three distinct symbols not in A , let $C = \{c, d, e\}$, and let $F = A \cup C$. For every TL_A -formula φ , set

$$\alpha_\varphi = c \wedge Xc \wedge X(CUGA) \wedge F(d \wedge X(cUd)) \wedge CU(A \wedge \varphi) ,$$

where A stands for $\bigvee_{a \in A} a$ and C for $c \vee d \vee e$.

The formula α_φ is chosen in such a way that for every $u \in F^\omega$ the following are equivalent:

- $u \models \alpha_\varphi$,
- u can be written as vw with $v \in C^+$ and $w \in A^\omega$ and such that $v \models c \wedge Xc$, $v \models F(d \wedge X(cUd))$, and $w \models \varphi$.

From [19] and [10], it follows that the set of finite words satisfying $c \wedge Xc$ and $F(d \wedge X(cUd))$ is not expressible in any of the fragments considered. So if φ is satisfiable, then α_φ is not expressible in any of the fragments. But if φ is not satisfiable, then so is α_φ , which means α_φ is expressible in any of the fragments considered. In other words, $\varphi \mapsto \alpha_\varphi$ is an appropriate reduction to prove the claim of the proposition. \square

Our upper bounds are as follows:

Theorem 9.3. *The $\{X, F\}$ -fragment is in E (exponential time), the other fragments listed in Table 1 are in $PSPACE$.*

Observe that the result for the $\{U\}$ -fragment is not new, but was already obtained in [20].

Proof. The proof is a refinement of the proof of Theorem 9.1.

Observe that each property expressed as forbidden pattern (as used in our characterizations) can not only be checked in polynomial time (which is folklore), it can also be checked non-deterministically in logarithmic space, simply by guessing the paths in questions, even if we are given a GCMA and need to check it on its left quotient. So if we interweave the construction of \mathbf{A}_φ , which has an exponential number of states, with the non-deterministic logarithmic-space tests for the existence of forbidden patterns, we obtain a polynomial-space procedure for testing the conditions on $T(\mathbf{A}_\varphi/\equiv_{\mathbf{A}_\varphi})$. (This is a standard argument in computational complexity.)

The situation is more complicated for the conditions on the loop languages. First observe that from the automaton \mathbf{A}_φ we can get reverse DFA's of size polynomial in the size of \mathbf{A}_φ such that every loop language is the union of the languages recognized by these reverse DFA's, which allows us to analyze the loop languages effectively.

We first deal with 1-local testability and stutter invariance and start with the observation that 1. and 2. from the proof of Theorem 9.1 can be adapted as follows. There are two states p and q in \mathbf{A}_φ that are not equivalent with respect to $\equiv_{\mathbf{A}}$ and such that one of the following conditions is true:

- (1) There are words $u, v \in A^*$ and there is a letter $a \in A$ such that $uav \in LL(p)$ and $uaav \in LL(q)$.
- (2) There are words $u, v \in A^*$ and letters $a, b \in A$ such that $uabv \in LL(p)$ and $ubav \in LL(q)$.

From this, it follows that we can bound the length of u and v polynomially in the size of \mathbf{A}_φ , which again yields polynomial-space procedures for both, 1-local testability and stutter invariance.

For (general) local testability, we apply the polynomial-time decision procedure for local testability developed in [16] to the product of the reverse DFA's mentioned above, which yields an exponential-time algorithm altogether. \square

We conclude this section with a more general version of Theorem 9.1:

Corollary 9.4. *For each of the fragments listed in Table 1, the following is decidable. Given an ω -regular language L , is L definable in the fragment?*

Proof. Given L we can construct effectively a Büchi automaton \mathbf{A} which recognizes L , see [28] for example. In [7] the decidability of the LTL-definability of $L(\mathbf{A}) = L$ is shown. Theorem 2.1 yields a GCMA \mathbf{B} with $L(\mathbf{B}) = L(\mathbf{A})$. Theorem 9.1 completes the proof. \square

10. OPEN PROBLEMS

We would like to state some questions:

- (1) Our lower and upper bounds for the complexity of the $\{X, F\}$ -fragment don't match. What is the exact complexity of this fragment?
- (2) Clearly, from our proofs it can be deduced that if a formula φ is equivalent to a formula in a fragment, an equivalent formula can be constructed effectively. What is the complexity of this construction task?
- (3) It is not difficult to come up with examples where every equivalent formula has exponential size (even exponential circuit size). What is the worst-case blow-up?— Observe that, in terms of circuit size, there is a polynomial upper bound for the $\{U\}$ -fragment, see [17].

REFERENCES

- [1] Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. M.I.T. Press, Boston, Mass., 2008.
- [2] Janusz A. Brzozowski and Imre Simon. Characterizations of locally testable events. *Discrete Math.*, 4(3):243 – 271, 1973.
- [3] Olivier Carton and Max Michel. Unambiguous Büchi automata. In Gaston H. Gonnet, Daniel Panario, and Alfredo Viola, ed., *LATIN*, vol. 1776 of *Lect. Notes in Comput. Sci.*, pages 407–416. Springer, 2000.
- [4] Olivier Carton and Max Michel. Unambiguous Büchi automata. *Theor. Comput. Sci.*, 297:37–81, 2003.
- [5] Joëlle Cohen, Dominique Prion, and Jean-Éric Pin. On the expressive power of temporal logic. *J. Comput. System Sci.*, 46(3):271–294, 1993.
- [6] Thomas Colcombet. Forms of determinism for automata (invited talk). In Christoph Dürr and Thomas Wilke, ed., *STACS*, vol. 14 of *LIPICs*, pages 1–23. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2012.
- [7] Volker Diekert and Paul Gastin. First-order definable languages. In Jörg Flum, Erich Grädel, and Thomas Wilke, ed., *Logic and Automata*, vol. 2 of *Texts in Logic and Games*, pages 261–306. Amsterdam University Press, 2008.
- [8] Volker Diekert and Manfred Kufleitner. Fragments of first-order logic over infinite words. *Theory Comput. Syst.*, 48(3):486–516, 2011.
- [9] E. Allen Emerson. Temporal and modal logic. In *Handbook of Theoretical Computer Science, Vol. B*, pages 995–1072. Elsevier, Amsterdam, 1990.
- [10] Kousha Etessami and Thomas Wilke. An until hierarchy and other applications of an Ehrenfeucht-Fraïssé game for temporal logic. *Inf. Comput.*, 160(1-2):88–108, 2000.
- [11] Dov M. Gabbay, Ian Hodkinson, and Mark Reynolds. *Temporal logic: Mathematical Foundations and Computational Aspects. Vol. 1*. Clarendon Press, New York, 1994.
- [12] Rob Gerth, Doron Peled, Moshe Y. Vardi, and Pierre Wolper. Simple on-the-fly automatic verification of linear temporal logic. In Piotr Dembinski and Marek Sredniawa, ed., *Protocol Specification, Testing and Verification*, vol. 38 of *IFIP Conference Proceedings*, pages 3–18. Chapman & Hall, 1995.
- [13] John E. Hopcroft, Rajeev Motwani, and Jeffrey D. Ullman. *Introduction to automata theory, languages, and computation*, 3rd edition. Addison-Wesley, Reading, Mass., 2006.
- [14] Leslie Lamport. The Temporal Logic of Actions. In *ACM Trans. Program. Lang. Syst.*, 16(3):872 – 923, 1994.
- [15] Hans Kamp. *Tense logic and the theory of linear order*. PhD thesis, University of California, Los Angeles, 1968.
- [16] Sam M. Kim, Robert McNaughton, and Robert McCloskey. A polynomial time algorithm for the local testability problem of deterministic finite automata. *IEEE Trans. Comput.*, 40:1087–1093, 1991.
- [17] Kousha Etessami. A note on a question of Peled and Wilke regarding stutter-invariant LTL. *Inform. Process. Lett.*, 75(6):261 – 263, 2000.
- [18] Robert McNaughton and Seymour A. Papert. *Counter-free automata*. M.I.T. Press, Boston, Mass., 1971.

- [19] Doron Peled and Thomas Wilke. Stutter-invariant temporal properties are expressible without the next-time operator. *Inform. Process. Lett.*, 63(5):243–246, 1997.
- [20] Doron Peled, Thomas Wilke, and Pierre Wolper. An algorithmic approach for checking closure properties of temporal logic specifications and ω -regular languages. *Theor. Comput. Sci.*, 195(2):183–203, 1998.
- [21] Dominique Perrin. Recent results on automata and infinite words. In Michal Chytil and Václav Koubek, ed., *MFCS*, vol. 176 of *Lect. Notes in Comput. Sci.*, pages 134–148. Springer, 1984.
- [22] Dominique Perrin and Jean-Éric Pin. *Infinite Words: Automata, Semigroups, Logic and Games*, vol. 141 of *Pure and Applied Mathematics*. Elsevier, Amsterdam, 2004.
- [23] Amir Pnueli. The temporal logic of programs. In *FOCS*, pages 46–57. IEEE, 1977.
- [24] Sebastian Preugschat and Thomas Wilke. Effective characterizations of simple fragments of temporal logic using prophetic automata. In Lars Birkedal, ed., *FoSSaCS*, vol. 7213 of *Lect. Notes in Comput. Sci.*, pages 135–149. Springer, 2012.
- [25] Marcel P. Schützenberger. On finite monoids having only trivial subgroups. *Inform. and Control*, 8(2):190 – 194, 1965.
- [26] A. Prasad Sistla and Edmund M. Clarke. The complexity of propositional linear temporal logics. *J. ACM*, 32:733–749, 1985.
- [27] Wolfgang Thomas. Star-free regular sets of ω -sequences. *Inform. and Control*, 42(2):148 – 156, 1979.
- [28] Wolfgang Thomas. Automata on infinite objects. In *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics (B)*, pages 133–192. 1990.
- [29] Moshe Y. Vardi and Pierre Wolper. Reasoning about infinite computations. *Inform. and Comput.*, 115(1):1–37, 1994.
- [30] Thomas Wilke. Classifying discrete temporal properties. Post-doctoral thesis, Christian-Albrechts-Universität zu Kiel, 1998.
- [31] Pierre Wolper, Moshe Y. Vardi, and A. Prasad Sistla. Reasoning about infinite computation paths (extended abstract). In *FOCS*, pages 185–194. IEEE, 1983.