

Anchored LTL Separation

Grgur Petric Maretić

Department of Computer Science
ETH Zürich
pgrgur@inf.ethz.ch

Mohammad Torabi Dashti

Department of Computer Science
ETH Zürich
torabidm@inf.ethz.ch

David Basin

Department of Computer Science
ETH Zürich
basin@inf.ethz.ch

Abstract

Gabbay’s separation theorem is a fundamental result for linear temporal logic (LTL). We show that separating a restricted class of LTL formulas, called anchored LTL, is elementary if and only if the translation from LTL to the linear temporal logic with only future temporal connectives is elementary. To prove this result, we define a canonical separation for LTL, and establish a correspondence between a canonical separation of anchored LTL formulas and the ω -automata that recognize these formulas.

The canonical separation of anchored LTL formulas has two further applications. First, we constructively prove that the safety closure of any LTL property is an LTL property, thus proving the decomposition theorem for LTL: every LTL formula is equivalent to the conjunction of a safety LTL formula and a liveness LTL formula. Second, we characterize safety, liveness, absolute liveness, stable, and fairness properties in LTL. Our characterization is effective: We reduce the problem of deciding whether an LTL formula defines any of these properties to the validity problem for LTL.

Categories and Subject Descriptors Theory of computation [Logic]: Modal and temporal logics; Theory of computation [Formal languages and automata theory]: Automata over infinite objects; Theory of computation [Formal languages and automata theory]: Regular languages

Keywords Linear Temporal Logic, ω -automata, Safety, Liveness

1. Introduction

The linear temporal logic with both past and future temporal connectives (LTL), and the linear temporal logic with only future temporal connectives (FLTL) are equally expressive [GPSS80]. This is because one can translate from LTL to FLTL using Gabbay’s separation algorithm [G87]. This may result in a nonelementary blow-up in the formula size.

The exact status of the blow-up, also referred to as the succinctness gap, is unknown. Markey [M03] has shown that LTL is at least exponentially more succinct than FLTL, i.e. there are arbitrarily large LTL formulas that cannot be expressed in FLTL without at least an exponential blow-up in the formula size. A tight bound on separation is also unknown. While it is expected that separation

is nonelementary, the succinctness gap is believed to be elementary [HR05]. The latter belief has become folklore, possibly due to a misinterpretation of Wilke’s construction of FLTL formulas from counter-free automata on *finite* words [W99]. Our main technical result connects the succinctness gap to the complexity of separating a strict subset of LTL formulas, called anchored formulas.

We use Gabbay’s separation theorem, stating that every LTL formula can be rewritten into a separated form, i.e. as a Boolean combination of future-only formulas and past-only formulas. We introduce a canonical form for separated formulas with a unique structure that simplifies our reasoning. We focus on the (canonical) separation of anchored LTL formulas. An anchored formula has the property that its semantics is the same at each moment in time. Intuitively, such a formula is bound to the initial time. For any given formula, we can obtain an equivalent anchored formula using simple syntactic manipulation.

Canonical separation of anchored formulas has two useful features. First, to translate an LTL formula φ into FLTL it suffices to separate its anchored equivalent. Second, separating anchored formulas appears to be easier than separating arbitrary formulas: We show that there is an elementary algorithm for separating anchored formulas if and only if there is an elementary algorithm for translating LTL to FLTL.

Canonical separation of anchored formulas is of further interest. We constructively prove that any property expressible in LTL is the intersection of a safety property and a liveness property, both expressible in LTL. This sharpens Alpern and Schneider’s safety-liveness decomposition [AS87] for ω -regular languages, and improves on our previous construction [PMTDB]. As a side result, we constructively show that for any safety LTL formula φ , there is a one-to-one correspondence between the *conjuncts* of any canonical separation of anchored φ and the states of a minimal deterministic Büchi automaton that recognizes the language of φ ; see Section 5.2. Furthermore, we characterize safety, liveness, absolute liveness, stable and fairness LTL formulas in terms of the separation of their anchored equivalents. The characterizations are effective, reducing the decision problem of recognizing such formulas to the validity problem for LTL.

Road Map. In Section 2 we recall linear temporal logic, safety and liveness, and ω -automata. In Section 3 we introduce canonical separation, and in Section 4 we relate the canonical separation of anchored formulas to Büchi automata, thereby relating the problem of separating anchored formulas to the problem of translating LTL to FLTL. In Section 5 we introduce an equivalence relation on the set of finite traces that generalizes the canonical separation of anchored formulas to arbitrary properties. Furthermore, we characterize safety and liveness in terms of this equivalence relation, constructively proving that the safety closure of every LTL formula is in LTL, and that LTL can be decomposed into safety and liveness within LTL. In Section 6 we characterize stable, absolute live-

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

CSL-LICS 2014, July 14–18, 2014, Vienna, Austria.
Copyright © 2014 ACM 978-1-4503-2886-9...\$15.00.
<http://dx.doi.org/10.1145/2603088.2603139>

ness, and fairness properties in LTL. In Section 7 we discuss related work.

2. Preliminaries

For a finite set AP of atomic propositions, we fix the alphabet $\Sigma = 2^{AP}$. Let Σ^ω be the set of all countably infinite sequences over Σ , Σ^+ be the set of all finite nonempty sequences over Σ , and $\Sigma^* = \Sigma^+ \cup \{\epsilon\}$, where ϵ is the empty sequence. An element of Σ^ω is a **path**, and a **property** is any set of paths. A **trace** is an element of Σ^+ . The **length** of a trace $t = p_0 p_1 \cdots p_i$, denoted $|t|$, is $i + 1$.

For a path $\pi = p_0 p_1 p_2 \cdots$, its **prefix** π_i is the trace $p_0 p_1 \cdots p_i$, and its **suffix** π^i is the path $p_i p_{i+1} \cdots$. For a trace t and a path (or a trace) π , the **concatenation** of t and π is denoted $t\pi$. The concatenation of ϵ and π is π .

2.1 Linear Temporal Logic

Our definitions in this section are standard; see for example [G87, M03, HR05].

Definition 1 (LTL Syntax). *The syntax of Linear Temporal Logic, LTL, is given by the grammar*

$$\varphi ::= \top \mid a \mid \neg\varphi \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \bullet\varphi \mid \circ\varphi \mid \varphi \mathcal{S}\psi \mid \varphi \mathcal{U}\psi,$$

where $a \in AP$. We write \perp for $\neg\top$, $\varphi \rightarrow \psi$ for $\neg\varphi \vee \psi$, $\varphi \leftrightarrow \psi$ for $(\varphi \rightarrow \psi) \wedge (\psi \rightarrow \varphi)$, $\blacklozenge\varphi$ for $\top \mathcal{S}\varphi$, $\blacksquare\varphi$ for $\neg\blacklozenge\neg\varphi$, $\diamond\varphi$ for $\top \mathcal{U}\varphi$, and $\square\varphi$ for $\neg\diamond\neg\varphi$. \triangle

The **size** of an LTL formula is defined inductively over its structure: $|\top| = 1$, $|a| = 1$ for $a \in AP$, $|\bullet\varphi| = 1 + |\varphi|$ for $\bullet \in \{\neg, \bullet, \circ\}$, and $|\varphi \star \psi| = 1 + |\varphi| + |\psi|$ for $\star \in \{\vee, \wedge, \mathcal{S}, \mathcal{U}\}$.

Definition 2 (LTL Semantics). *For a path $\pi = p_0 p_1 p_2 \cdots$ and $i \in \mathbb{N}$, the **satisfaction relation** for LTL formulas is defined inductively over the formula structure:*

$$\begin{aligned} \pi, i &\models \top \\ \pi, i &\models a && \text{if } a \in p_i \\ \pi, i &\models \neg\varphi && \text{if } \pi, i \not\models \varphi \\ \pi, i &\models \varphi \vee \psi && \text{if } \pi, i \models \varphi \text{ or } \pi, i \models \psi \\ \pi, i &\models \varphi \wedge \psi && \text{if } \pi, i \models \varphi \text{ and } \pi, i \models \psi \\ \pi, i &\models \bullet\varphi && \text{if } i > 0 \text{ and } \pi, i-1 \models \varphi \\ \pi, i &\models \circ\varphi && \text{if } \pi, i+1 \models \varphi \\ \pi, i &\models \varphi \mathcal{S}\psi && \text{if there is a } j \leq i \text{ such that } \pi, j \models \psi \\ &&& \text{and } \pi, k \models \varphi \text{ for all } j < k \leq i \\ \pi, i &\models \varphi \mathcal{U}\psi && \text{if there is a } j \geq i \text{ such that } \pi, j \models \psi \\ &&& \text{and } \pi, k \models \varphi \text{ for all } i \leq k < j \end{aligned}$$

When $\pi, i \models \varphi$, we say π **satisfies** φ at **time** i and we say φ is **satisfiable** if there is a path π such that $\pi, 0 \models \varphi$. The formula φ is **valid** if $\neg\varphi$ is not satisfiable. An LTL formula φ **expresses** the property $L(\varphi) = \{\pi \in \Sigma^\omega \mid \pi, 0 \models \varphi\}$. \triangle

A **past** LTL formula is a formula that does not use the future temporal connectives \circ and \mathcal{U} . A **future** LTL formula is a formula that does not use the past temporal connectives \bullet and \mathcal{S} .

Lemma 3. *Let F be a future formula and let P be a past formula. Then for every $i \in \mathbb{N}$ and every path $\pi \in \Sigma^\omega$*

1. $\pi, i \models F$ iff $t\pi^i, |t| \models F$, for every trace $t \in \Sigma^+$, and
2. $\pi, i \models P$ iff $\pi_i\sigma, i \models P$, for every path $\sigma \in \Sigma^\omega$. \blacktriangle

Informally, the satisfaction of future formulas is independent of changing the past and the satisfaction of past formulas is independent of changing the future. This justifies the following definition of satisfaction for past formulas on finite traces. For a past formula P and a trace $t \in \Sigma^+$, we write $t \models P$ if $t\pi, |t| - 1 \models P$, for any path π .

Definition 4. *We define two equivalence relations on LTL formulas.*

1. The formulas φ and ψ are **globally equivalent**, denoted $\varphi \equiv_g \psi$, if $\forall \pi \in \Sigma^\omega. \forall i \in \mathbb{N}. (\pi, i \models \varphi \iff \pi, i \models \psi)$.
2. The formulas φ and ψ are **initially equivalent**, denoted $\varphi \equiv \psi$, if $\forall \pi \in \Sigma^\omega. (\pi, 0 \models \varphi \iff \pi, 0 \models \psi)$. \triangle

Global equivalence implies initial equivalence, but the converse is false. For example, the formulas $\bullet\top$ and \perp are initially equivalent, but they are not globally equivalent. Note that $\varphi \equiv \psi$ if and only if $L(\varphi) = L(\psi)$. Checking the global equivalence of φ and ψ can be reduced to the validity problem for LTL. Namely, $\varphi \equiv_g \psi$ if and only if the formula $\square(\varphi \leftrightarrow \psi)$ is valid.

The separation theorem of Gabbay [G87] states that every LTL formula is globally equivalent to a Boolean combination of past and future formulas. Furthermore, Gabbay gives an algorithm for separation. It is easy to rewrite Gabbay's Boolean combination as a conjunction of implications as stated in the following theorem.

Theorem 5 (LTL Separation). *Each LTL formula φ is globally equivalent, through a sequence of rewrites, to a formula*

$$(P_1 \rightarrow \circ F_1) \wedge \cdots \wedge (P_n \rightarrow \circ F_n),$$

where P_i are past formulas and F_i are future formulas. \blacktriangle

We call this formula a **separation** of φ . We write $(P \rightarrow \circ F)_n$ to denote a separation consisting of n conjuncts. Each **conjunct** is an implication $P_i \rightarrow \circ F_i$, and we call the formulas P_i and F_i the past part and the future part of the conjunct, respectively.

We remark that Gabbay's algorithm results in a **nonelementary** blow-up: there is no constant $k \in \mathbb{N}$ such that for any formula φ the size of its separation is bounded from above by a k -story exponential function in the size of φ ; for a definition of nonelementary complexity see [M74]. We refer to the minimum blow-up in the formula size that is introduced by separation as the **separation gap**. The exact status of the separation gap is an open problem [HR05].

Example 6. *We use the formula $\varphi = \square(a \rightarrow \blacklozenge b)$ as our running example. This formula expresses a simple precedence property, for example that access to some resource (a) must be preceded by a request (b). A separation of φ is given by the formula*

$$\psi = ((\blacksquare\neg b \wedge a) \rightarrow \circ \perp) \wedge (\blacksquare\neg b \rightarrow \circ(\square\neg a \vee \neg a \mathcal{U} b)).$$

Intuitively, with no requests in the past, access is prohibited in the present, and a request will be made in the future before the first access, or no access will ever be made in the future. \triangle

2.2 Safety and Liveness

According to Lamport [L77], a safety property states that something (usually "bad") never happens. For example, the elevator door never opens while the elevator is moving. A liveness property states that something (usually "good") eventually happens. For example, every process in a multitasking system will eventually be granted CPU time. There are various formalizations of safety and liveness, such as [LPZ85], but Alpern and Schneider's formalization [AS85] has become widely accepted.

Definition 7 ([AS85]). *A set S of paths is a **safety property** if for every $\pi \notin S$ there is an $i \in \mathbb{N}$ such that $\pi_i\sigma \notin S$ for all paths $\sigma \in \Sigma^\omega$.*

*A set L of paths is a **liveness property** if for every $t \in \Sigma^*$ there is a path π such that $t\pi \in L$.* \triangle

Intuitively, every path that is not in a safety property S has a finite "bad" prefix, while every trace has a "good" suffix for a liveness property L . For an LTL formula φ , we say φ is a liveness, respectively safety, formula if $L(\varphi)$ is a liveness, respectively safety, property. From this point on, whenever confusion is unlikely, we identify a formula with the property it expresses.

Every property can be expressed as the intersection of a safety property and a liveness property [AS85]. Moreover, for every property \mathcal{L} we can define its **safety closure** $[\mathcal{L}]$ as the smallest safety

property containing \mathcal{L} . In [PMTDB], we prove that for every property expressible in LTL, its safety closure is also expressible in LTL. As a consequence, the decomposition into safety and liveness is possible within LTL.

Theorem 8 ([PMTDB]). *For every LTL formula φ , the safety closure $\lceil L(\varphi) \rceil$ of $L(\varphi)$ is expressible by an LTL formula. \blacktriangle*

Theorem 9 (Safety-Liveness Decomposition in LTL [PMTDB]). *For every LTL formula φ , there are LTL formulas σ and λ such that $L(\sigma)$ is safety, $L(\lambda)$ is liveness, and $\varphi \equiv \sigma \wedge \lambda$. \blacktriangle*

In Section 5.2 we use canonical separation to constructively prove these theorems and to characterize safety and liveness in LTL. See Section 7 for our comparison with [PMTDB].

2.3 ω -automata

We assume that the reader is familiar with automata on finite words and regular expressions. For background see [H79, M71]. We take the definition of Büchi automata from [T90].

Definition 10. A Büchi automaton over the alphabet Σ is a tuple $\mathcal{A} = (Q, q_0, \Delta, F)$ with a finite set Q of states, an initial state $q_0 \in Q$, a transition relation $\Delta \subseteq Q \times \Sigma \times Q$, and a set $F \subseteq Q$ of accepting states. A Büchi automaton is **deterministic** if $\Delta: Q \times \Sigma \rightarrow Q$ is a partial transition function.

A **run** of \mathcal{A} on a path $\pi = p_0 p_1 \dots \in \Sigma^\omega$ is a sequence of states $s_0 s_1 \dots$ such that $s_0 = q_0$ and $(s_i, p_i, s_{i+1}) \in \Delta$, for $i \in \mathbb{N}$. The run is **accepting** if some state of F occurs infinitely often in the run. \mathcal{A} **accepts** π if there is an accepting run of \mathcal{A} on π . The property $L(\mathcal{A}) = \{\pi \in \Sigma^\omega \mid \mathcal{A} \text{ accepts } \pi\}$ is the property recognized by \mathcal{A} .

For $t \in \Sigma^+$ we define $(s, t, s') \in \Delta^+$ if there is a sequence of states $s_0, s_1, \dots, s_{|t|}$ such that $s = s_0, s' = s_{|t|}$ and $(s_i, t_i, s_{i+1}) \in \Delta$, for $0 \leq i < |t|$. Δ

For a non-deterministic Büchi automaton, we allow a set of initial states $Q_0 \subseteq Q$ instead of a single state q_0 . Formally, the notation $\mathcal{A} = (Q, Q_0, \Delta, F)$ is a short-hand for the automaton $\mathcal{A}' = (Q \cup \{q_0\}, q_0, \Delta', F)$, where $q_0 \notin Q$ and the relation Δ' is defined as follows. For all states $p, q \in Q \cup \{q_0\}$ and every $a \in \Sigma$, the tuple (p, a, q) is in Δ' if $(p, a, q) \in \Delta$, or $p = q_0$ and there is a state $s \in Q_0$ such that $(s, a, q) \in \Delta$.

Several types of ω -automata have been studied in the literature. These include generalized Büchi automata; see for example [GO01]. They differ from Büchi automata only in their acceptance condition. In a generalized Büchi automaton, we have $F \subseteq 2^Q$ and a run is accepting if it visits at least one element from each set in F infinitely often.

3. Canonical Separation

We start with an informal account of canonical separation and anchored formulas and provide the definitions afterward.

Consider the following scenario. A system that has been running for some time is required to satisfy a property expressed by an LTL formula φ (at the initial moment in time). The past behavior, logged as a trace t , determines which future system behavior is allowed and which would falsify φ . Intuitively, a separated formula $\psi = (P_1 \rightarrow \bigcirc F_1) \wedge \dots \wedge (P_n \rightarrow \bigcirc F_n)$ has a suitable structure for this scenario: the past formulas P_i stand for the possible pasts, while the corresponding future formulas F_i describe the allowed futures. Specifically, our goal is to produce a separated formula ψ such that for every trace t there is a unique index i such that $t \models P_i$ (determines the past) and that for every path π the path $t\pi$ satisfies φ if and only if $\pi, 0 \models F_i$. We will show that $\psi \not\equiv_g \varphi$, but ψ can be constructed from φ . We require that the separation ψ meets the following conditions.

- (1) For every $t \in \Sigma^+$, there is a unique i such that $t \models P_i$.
- (2) For every $\pi \in \Sigma^\omega$, if there is a $k \in \mathbb{N}$ such that $\pi, k \models \varphi$ then $\pi, 0 \models \varphi$.
- (3) For every $\pi \in \Sigma^\omega$, if $\pi, 0 \models \varphi$ then $\pi, k \models \psi$, for any $k \in \mathbb{N}$.
- (4) For all i, j , if $i \neq j$ then $F_i \not\equiv F_j$.

The intuition behind these conditions is as follows. The first condition guarantees that ψ defines a partition on the set of finite traces. From the second condition, it follows that if $t \models P_i$ then for every path π that satisfies F_i , the path $t\pi$ satisfies φ initially. The third condition ensures that if $t \models P_i$ and π falsifies F_i , the path $t\pi$ falsifies φ initially. The last condition requires that we group together all traces that allow exactly the same future behaviors.

In general, a separation of the formula φ need not satisfy any of these conditions. We first define anchored formulas and prove that anchored φ satisfies (2) and (3) for any formula φ . Note that any separation of anchored φ , being globally equivalent to anchored φ , therefore also satisfies (2) and (3). We then define the notion of a canonical separation to satisfy the remaining two conditions for ψ .

Definition 11. We say an LTL formula ψ is **anchored** if, for every path π , either π satisfies ψ at every time i or π falsifies ψ at every time i . For an LTL formula φ , we refer to the formula $\blacklozenge\varphi$ as **anchored φ** . Δ

Note that anchored LTL formulas should not be confused with Manna and Pnueli's "anchored temporal framework" [MP89].

It is immediate that anchored φ is initially equivalent to φ . That anchored φ is indeed anchored is entailed by the following lemma.

Lemma 12. *For every LTL formula φ , anchored φ satisfies conditions (2) and (3).*

Proof. Let $\pi \in \Sigma^\omega$ and $i \in \mathbb{N}$ be such that $\pi, i \models \blacklozenge\varphi$. In particular, $\pi, 0 \models \blacklozenge\varphi$ and therefore $\pi, 0 \models \varphi$, proving (2).

Let $\pi \in \Sigma^\omega$ such that $\pi, 0 \models \varphi$. Then, $\pi, j \models \blacklozenge\varphi$ for every $j \in \mathbb{N}$. In particular, for every $i \in \mathbb{N}$ whenever $j \leq i$ then also $\pi, j \models \blacklozenge\varphi$. Thus, $\pi, i \models \blacklozenge\varphi$ for every i , proving (3). \blacktriangle

We now define a canonical separation that satisfies the remaining two conditions for ψ . Intuitively, it is a separation that satisfies (1) and (4); moreover, it has a distinguished conjunct with an unsatisfiable *future* part, and the *past* part of all other conjuncts is satisfiable. We then constructively prove that such a formula exists, that it can be constructed from any separation, and that all such formulas have the same structure, justifying the name "canonical".

Definition 13 (Canonical Separation). *For an LTL formula ψ , a **canonical separation** of ψ is a separation $(P \rightarrow \bigcirc F)_n$ that satisfies the following properties.*

- For every $t \in \Sigma^+$, there is a unique i such that $t \models P_i$.
- For all i, j , if $i \neq j$ then $F_i \not\equiv F_j$.
- There is an i such that $F_i \equiv \perp$.
- If there is an i such that $P_i \equiv_g \perp$, then $F_i \equiv \perp$. Δ

We show that any separation can be transformed into a canonical one. Below, we write $[1, n]$ for the set $\{1, \dots, n\}$.

Construction 14. *Let $(P \rightarrow \bigcirc F)_n$ be any separation of an LTL formula ψ . Consider the formula*

$$\bigwedge_{I \subseteq [1, n]} \left(\left(\left(\bigwedge_{i \in I} P_i \right) \wedge \left(\bigwedge_{j \in [1, n] \setminus I} \neg P_j \right) \right) \rightarrow \bigcirc \bigwedge_{i \in I} F_i \right),$$

where $\bigwedge_{i \in \emptyset} = \top$. We refer to the past and future parts of the conjuncts of this formula by P'_i and F'_i , respectively. Note that the past parts are pairwise inconsistent. It is immediate that this formula is

also a separation of ψ . We apply the following transformations, in order.

- (i) For every i , if $P'_i \equiv_g \perp$, eliminate the conjunct.
- (ii) Add the conjunct $(\perp \rightarrow \bigcirc \perp)$.
- (iii) For every i , if there is a $j \neq i$ such that $F'_i \equiv F'_j$, merge both conjuncts to a single conjunct $(P'_i \vee P'_j \rightarrow F'_i)$.

It is immediate that this process terminates and that the resulting formula is a canonical separation of ψ . \triangle

Since every separation of any LTL formula is globally equivalent to the formula, it follows from Lemma 12 that any canonical separation of anchored φ satisfies conditions (2) and (3). Therefore, for any formula φ , every canonical separation of anchored φ satisfies conditions (1)–(4).

Example 15. We construct a canonical separation of anchored φ for the precedence property $\varphi = \Box(a \rightarrow \blacklozenge b)$. A separation of anchored φ is given by the formula $(P_1 \rightarrow \bigcirc F_1) \wedge (P_2 \rightarrow \bigcirc F_2)$, where $P_1 = \blacksquare \neg b$, $F_1 = (\neg a \mathcal{U} b) \vee \Box \neg a$, $P_2 = \neg \blacksquare (a \rightarrow \blacklozenge b)$, and $F_2 = \perp$. We apply Construction 14 to transform this separation into a canonical separation of anchored φ , which is the conjunction of the following three formulas:

$$\begin{aligned} (\blacksquare (a \rightarrow \blacklozenge b) \wedge \blacklozenge b) &\rightarrow \bigcirc \top, \\ \blacksquare (\neg a \wedge \neg b) &\rightarrow \bigcirc ((\neg a \mathcal{U} b) \vee \Box \neg a), \\ \blacklozenge (a \wedge \blacksquare \neg b) &\rightarrow \bigcirc \perp. \end{aligned} \quad \triangle$$

We are only interested in the semantics of a canonical separation's conjuncts, and not in their particular syntax. The following theorem states that canonical separations are canonical in the sense that for any formula ψ , all canonical separations of ψ have the same structure. More precisely, we show that there is no semantic distinction between corresponding conjuncts of two canonical separations of the same formula.

Theorem 16. Let $(P \rightarrow \bigcirc F)_n$, and $(P' \rightarrow \bigcirc F')_m$ be two canonical separations of the LTL formula ψ . Then, they have the same structure, namely

1. $n = m$, and
2. for every i there is a j such that $P_i \equiv_g P'_j$ and $F_i \equiv F'_j$.

Proof. Let $P \rightarrow \bigcirc F$ be a conjunct from the first canonical separation such that P is satisfiable. Take any trace t such that $t \models P$. By the definition of canonical separation, there is a conjunct $P' \rightarrow \bigcirc F'$ in the second canonical separation such that $t \models P'$. It follows that for every $\pi \in \Sigma^\omega$, the satisfaction $\pi, |t| - 1 \models \psi$ is equivalent to $\pi, 0 \models F$ and to $\pi, 0 \models F'$. Therefore, F and F' are initially equivalent. Assume $P \not\equiv_g P'$. Without loss of generality, let $t' \models P'$ and $t' \not\models P$. Then there is another conjunct, $P'' \rightarrow \bigcirc F''$, in the first canonical separation such that $t' \models P''$. Now, it follows that $F' \equiv F''$, but then $F'' \equiv F$, which contradicts the definition of a canonical separation.

In case one of the canonical separations has a conjunct $P \rightarrow \bigcirc F$ such that P is satisfiable and $F \equiv \perp$, it follows that the other separation has such a conjunct as well. Otherwise, both formulas must have a conjunct $\perp \rightarrow \bigcirc \perp$. This gives us a one-to-one correspondence between the conjuncts of both formulas, and thus $n = m$. \blacktriangle

To simplify discussions, for any canonical separation, we will use the distinguished index \perp to refer to the conjunct $(P_\perp \rightarrow \bigcirc F_\perp)$, where $F_\perp \equiv \perp$.

4. Connection to ω -Automata

In this section, we establish a connection between canonical separation of anchored formulas and ω -automata. This will allow us to

give upper bounds on the number of conjuncts and the size of the past parts in the separation. Furthermore, we show that separating anchored formulas is elementary if and only if the translation from LTL to FLTL is elementary.

Given a trace t and a formula φ , a canonical separation $(P \rightarrow \bigcirc F)_n$ of anchored φ can be used to determine the necessary and sufficient condition for a continuation π such that $t\pi$ satisfies φ . We simply find an index i such that $t \models P_i$ and the condition is given by F_i . Let us consider a Büchi automaton $\mathcal{A} = (Q, q_0, \Delta, F)$ recognizing $L(\varphi)$, and let $Q_t = \{s' \in Q \mid (q_0, t, s') \in \Delta^+\}$. The set Q_t consists of all the end states of finite runs of \mathcal{A} on t . Let $\mathcal{A}' = \mathcal{A}[q_0 \leftarrow Q_t]$ be the automaton (Q, Q_t, Δ, F) , in which the initial state q_0 is replaced by the set of states Q_t . We show that Q_t characterizes a subset of the set of all traces t' such that $t' \models P_i$, while \mathcal{A}' recognizes exactly the paths π such that $\pi, 0 \models F_i$.

Theorem 17. Let \mathcal{A} be a Büchi automaton and let t be a trace. A path π is accepted by the automaton \mathcal{A}' if and only if $t\pi \in L(\mathcal{A})$.

Proof. Assume there is an accepting run $s_0 s_1 \dots$ of \mathcal{A} on $t\pi$. It is immediate that $s_{|t|} \in Q_t$, and therefore $s_{|t|} s_{|t|+1} \dots$ is an accepting run of \mathcal{A}' on π . Conversely, let $\pi \in L(\mathcal{A}')$, and let $s_0 s_1 \dots$ be an accepting run on π . By definition, $s_0 \in Q_t$, and therefore there is a finite run $p_0 p_1 \dots p_{|t|}$ of \mathcal{A} on t such that $p_{|t|} = s_0$. It follows that $p_0 p_1 \dots p_{|t|} s_1 s_2 \dots$ is an accepting run of \mathcal{A} on $t\pi$. \blacktriangle

This establishes a connection between Büchi automata and a canonical separation of anchored φ , as stated in the following corollaries.

Corollary 18. Let φ be an LTL formula and let $(P \rightarrow \bigcirc F)_n$ be a canonical separation of anchored φ . For any automaton \mathcal{A} recognizing φ , the following holds.

$$\forall t \in \Sigma^+. \forall i \in [1, n]. t \models P_i \implies L(\mathcal{A}') = L(F_i). \quad \blacktriangle$$

Corollary 19. Let φ be an LTL formula and let $(P \rightarrow \bigcirc F)_n$ be a canonical separation of anchored φ . For any automaton \mathcal{A} recognizing φ and any two traces $t, t' \in \Sigma^+$ the following holds.

$$L(\mathcal{A}') = L(\mathcal{A}'') \implies \forall i \in [1, n]. (t \models P_i \iff t' \models P_i). \quad \blacktriangle$$

Note that Theorem 17 and its consequences translate easily to other types of ω -automata, in particular generalized Büchi automata. The proof is almost identical to the proof of Theorem 17.

The following example further explains the connection between canonical separations of anchored LTL formulas and their corresponding ω -automata. We assume that the reader is familiar with ω -regular expressions. We follow the usual precedence conventions: concatenation has highest priority, followed by the Kleene star ($*$) and cross (\cdot), ω -closure ($^\omega$), complement ($\bar{\cdot}$), and finally set union ($\cdot + \cdot$). For simplicity we use propositional formulas as syntactic sugar, where a formula Φ stands for the union of all literals $l \in \Sigma = 2^{AP}$ that represent a truth assignment satisfying Φ . For example, if $AP = \{a, b\}$, the formula $\Phi = a$ stands for the regular expression $(\{a\} + \{a, b\})$.


Example 20. Let φ be an LTL formula and let $(P \rightarrow \bigcirc F)_n$ be a canonical separation of anchored φ . For any automaton \mathcal{A} recognizing φ , the following holds.

$$\forall t, t' \in \Sigma^+. Q_t = Q_{t'} \implies \forall i \in [1, n]. (t \models P_i \iff t' \models P_i).$$

This statement is a special case of Corollary 19. The statement's converse is however false. In fact, there are formulas for which one cannot construct an automaton such that all traces that satisfy the same past part reach exactly the same set of states. Let φ be the formula

$$\Box \left(\diamond (\neg a \wedge b) \wedge \diamond (\neg a \wedge \neg b) \wedge \diamond (a \wedge \neg b) \wedge \diamond (a \wedge b) \right).$$

Figure 1. Automaton recognizing the fairness formula of Example 20



Intuitively, φ requires that all four combinations of truth values of a and b occur infinitely often. A canonical separation of anchored φ is given by $(\top \rightarrow \circ \varphi) \wedge (\perp \rightarrow \circ \perp)$. It is immediate that all finite traces satisfy the same past part, namely \top .

A Büchi automaton with the least number of states recognizing φ is given in Fig. 1. Even though all traces are equivalent with respect to the satisfaction of φ , the automaton has one non-accepting state for each combination of truth values of a and b , as well as an additional accepting state that serves to monitor that the four combinations all occur infinitely often. \triangle

While the separation gap is believed to be nonelementary [HR05], the connection between a canonical separation and ω -automata allows us to prove elementary bounds on the number of conjuncts and the size of the past parts of a canonical separation of anchored LTL formulas. Furthermore, we show that the size of the future parts is related to the succinctness gap between LTL and FLTL, which is believed to be elementary [HR05].

Let φ be an LTL formula and let \mathcal{A}_φ be a non-deterministic ω -automaton with m states that recognizes $L(\varphi)$. It follows from Corollary 18 that for any canonical separation $(P \rightarrow \circ F)_n$ of anchored φ , $n \leq 2^m$. Gastin and Oddoux [GO01] show that for any LTL formula φ , one can construct a generalized Büchi automaton \mathcal{A}_φ with at most $2^{|\varphi|+1}$ states. The following theorem is now immediate.

Theorem 21. *For any LTL formula φ , the number of conjuncts of a canonical separation of anchored φ is at most double exponential in the size of φ .* \blacktriangle

We have thus established an elementary upper bound on the number of conjuncts of canonical separations of anchored LTL formulas. Now, we turn to the size of the past parts.

Theorem 22. *Let φ be an LTL formula. There exists a canonical separation $(P \rightarrow \circ F)_n$ of anchored φ where the size of each past part P_i is at most quadruple exponential in the size of φ .*

Proof. Let $\mathcal{A}_\varphi = (Q, q_0, \Delta, F)$ be a generalized Büchi automaton recognizing $L(\varphi)$. We consider this automaton as a non-deterministic finite state automaton on finite words (NFA) with no accepting states. More precisely, we consider the NFA $\mathcal{A}_F = (Q, q_0, \Delta, \emptyset)$. We then determinize the NFA, using the standard subset construction, and obtain a deterministic finite automaton on finite words (DFA) $\mathcal{D} = (2^Q, \{q_0\}, \Delta', \emptyset)$. The key to the proof is that for each past part P_i of the separation, there is a set $\mathcal{F}^i \subseteq \mathcal{D}$

such that the DFA $\mathcal{D}^i = (2^Q, \{q_0\}, \Delta', \mathcal{F}^i)$ recognizes $L(P_i)$. This implies that \mathcal{D}^i is counter-free. We may then use the construction from [MP90, W99] to construct the formula P_i . The size of this formula is double exponential in the number of states of \mathcal{D}^i , whose size is in turn double exponential in the size of φ .

To construct \mathcal{F}^i , take any state $S \in \mathcal{D}$ and find all $S' \in \mathcal{D}$ such that $\mathcal{A}_\varphi[q_0 \leftarrow S] = \mathcal{A}_\varphi[q_0 \leftarrow S']$. By Corollary 19, any trace t reaching S and any trace t' reaching S' satisfy the past part of the same conjunct, say P_i . \blacktriangle

We are ready to prove our main theorem.

Theorem 23. *The following statements are logically equivalent:*

- (a) *The separation gap for anchored LTL is elementary.*
- (b) *The succinctness gap between LTL and FLTL is elementary.*

Proof. One direction is immediate. Take an arbitrary LTL formula φ . First, we anchor and then separate φ . Second, in every past formula in the separation, we iteratively replace $\psi_1 S \psi_2$ by ψ_2 and $\bullet \psi$ by \perp . This results in an FLTL formula, denoted φ^F , that is initially equivalent to φ . The size of φ^F is clearly smaller than the size of the separation. Therefore, assuming that the separation gap for anchored LTL formulas is elementary, there is an elementary bound on the succinctness gap between LTL and FLTL.

To prove the converse, we show that the size of the future parts is elementary in the size of φ^F . This, together with Theorems 21 and 22, proves our claim. So, assume that there is an elementary bound on the size of φ^F . For every index i and trace $t = t_0 t_1 \cdots t_{|t|-1}$ such that $t \models P_i$ we know that $t\pi, 0 \models \varphi^F$ if and only if $\pi, 0 \models F_i$. Intuitively, we will “partially evaluate” the satisfaction of φ^F on the trace t to obtain F_i . The procedure is similar to semantic tableaux for LTL. Every formula of the form $\psi_1 \mathcal{U} \psi_2$ is equivalent to the formula $\psi_2 \vee (\psi_1 \wedge \circ (\psi_1 \mathcal{U} \psi_2))$. Using this equivalence, we can rewrite φ^F to a formula in which only propositional formulas are not under the scope of the next time connective \circ . Since φ^F contains no past temporal connectives, the satisfaction of all formulas under the scope of the next time connective is independent of t_0 . We can therefore replace each formula a that is not under the scope of \circ with \top if $t_0 \models a$, and with \perp if $t_0 \not\models a$. We then rewrite the resulting formula to $\circ \varphi_1^F$. It is immediate that $t\pi, 0 \models \varphi^F$ iff $t_1 t_2 \cdots t_{|t|-1} \pi, 0 \models \varphi_1^F$. We continue this procedure for the entire t , and ultimately get $t\pi, 0 \models \varphi^F$ iff $\pi, 0 \models \varphi_{|t|}^F$. Therefore $\varphi_{|t|}^F \equiv F_i$. Note that, due to Corollary 19, one can always choose a trace t such that $t \models P_i$ and the length of t is at most double exponential in the size of φ .

To conclude the proof, we show that the size of $\varphi_{|t|}^F$ is at most exponential in $|\varphi^F|$. For any formula ψ , we define the **set of sub-formulas** of ψ , denoted $\text{SF}(\psi)$, inductively as follows. $\text{SF}(\top) = \{\top\}$, $\text{SF}(a) = \{a\}$ for $a \in AP$, $\text{SF}(\bullet \psi) = \{\bullet \psi\} \cup \text{SF}(\psi)$ for $\bullet \in \{\neg, \bullet, \circ\}$, and $\text{SF}(\psi_1 \star \psi_2) = \{\psi_1 \star \psi_2\} \cup \text{SF}(\psi_1) \cup \text{SF}(\psi_2)$ for $\star \in \{\vee, \wedge, \mathcal{S}, \mathcal{U}\}$. It is immediate that the number of elements of $\text{SF}(\varphi^F)$ is at most $|\varphi^F|$ and that each formula $\varphi_{|t|}^F$ is a Boolean combination of formulas in $\text{SF}(\varphi^F)$. In particular, $\varphi_{|t|}^F$ is a Boolean combination of formulas in $\text{SF}(\varphi^F)$. It follows that we can write $\varphi_{|t|}^F$ using Boolean connectives and at most $2^{|\varphi^F|}$ occurrences of sub-formulas of φ^F , each of size at most $|\varphi^F|$. \blacktriangle

5. Safety-Liveness Decomposition

In Section 3, we formalized how finite traces determine future behavior with respect to properties expressed by LTL formulas. Here, we generalize this idea to arbitrary properties. We show in Section 5.1 that, for any property \mathcal{L} , this reasoning can be formalized as an equivalence relation, denoted $\approx_{\mathcal{L}}$. Furthermore, the safety closure of \mathcal{L} and the condition for \mathcal{L} to be a liveness

property can be characterized in terms of equivalence classes of $\approx_{\mathcal{L}}$. In Section 5.2, we turn to safety-liveness decomposition in LTL.

5.1 Decomposition of Arbitrary Properties

For a property \mathcal{L} we define an equivalence relation $\approx_{\mathcal{L}}$ that relates two finite traces u and v if for every infinite continuation π either both $u\pi$ and $v\pi$ satisfy the property (i.e. both belong to \mathcal{L}), or both falsify it. Intuitively, two finite traces are equivalent if they place the same restrictions on future behaviors.

Definition 24. Let \mathcal{L} be a property. We define the equivalence relation $\approx_{\mathcal{L}} \subseteq \Sigma^+ \times \Sigma^+$ as follows.

$$u \approx_{\mathcal{L}} v \text{ if } \forall \pi \in \Sigma^\omega. (u\pi \in \mathcal{L} \iff v\pi \in \mathcal{L}). \quad \Delta$$

Definition 25 ([N58]). An equivalence relation \approx on a set of traces Σ^+ is **right invariant** with respect to concatenation if for all traces $u, v, z \in \Sigma^+$, $u \approx v$ implies $uz \approx vz$. Δ

It is immediate that $\approx_{\mathcal{L}}$ is right invariant. For a property \mathcal{L} , if there is a trace $v \in \Sigma^+$ such that $v\pi \notin \mathcal{L}$ for every $\pi \in \Sigma^\omega$, we define the set \mathcal{L}_\perp of irremediable traces as the equivalence class $[v]_{\approx_{\mathcal{L}}}$. Otherwise, we define $\mathcal{L}_\perp = \emptyset$. The safety closure of \mathcal{L} and the condition for \mathcal{L} to be a liveness property can be characterized using \mathcal{L}_\perp .

Lemma 26. For every property \mathcal{L} , its safety closure $\lceil \mathcal{L} \rceil$ is the set $\{\pi \in \Sigma^\omega \mid \forall i \in \mathbb{N}. \pi_i \notin \mathcal{L}_\perp\}$.

Proof. We show that $S = \{\pi \in \Sigma^\omega \mid \forall i \in \mathbb{N}. \pi_i \notin \mathcal{L}_\perp\}$ is a safety property containing \mathcal{L} , and that it is a subset of every other such property. It follows that $S = \lceil \mathcal{L} \rceil$.

For any path $\pi \notin S$, there is an $i \in \mathbb{N}$ such that $\pi_i \in \mathcal{L}_\perp$. It follows that $\pi_i\sigma \notin S$ for every path σ and therefore S is a safety property. Since $\pi_i \in \mathcal{L}_\perp$, in particular $\pi \notin \mathcal{L}$ and thus $\mathcal{L} \subseteq S$. Finally, let S' be a safety property containing \mathcal{L} . Then, for every path $\pi \notin S'$, there is an $i \in \mathbb{N}$ such that $\pi_i\sigma \notin S'$, for every path σ . Since $\mathcal{L} \subseteq S'$, also $\pi_i\sigma \notin \mathcal{L}$ for every path σ . Therefore, $\pi_i \in \mathcal{L}_\perp$ and thus $\pi \notin S$. It follows that $S \subseteq S'$. \blacktriangle

The following proposition shows that for any property \mathcal{L} , the relation $\approx_{\mathcal{L}}$ is a refinement of $\approx_{\lceil \mathcal{L} \rceil}$.

Proposition 27. For any property \mathcal{L} , for any $u, v \in \Sigma^+$ if $u \approx_{\mathcal{L}} v$, then $u \approx_{\lceil \mathcal{L} \rceil} v$. Furthermore, $\mathcal{L}_\perp = \lceil \mathcal{L} \rceil_\perp$.

Proof. Let $u \approx_{\mathcal{L}} v$, and let $\pi \in \Sigma^\omega$ be an arbitrary path. Since $\approx_{\mathcal{L}}$ is right invariant, it follows that for every $i \in \mathbb{N}$ also $u\pi_i \approx_{\mathcal{L}} v\pi_i$. Therefore, there is an i such that $u\pi_i \in \mathcal{L}_\perp$ iff there is an i such that $v\pi_i \in \mathcal{L}_\perp$ and thus $u\pi \in \lceil \mathcal{L} \rceil$ iff $v\pi \in \lceil \mathcal{L} \rceil$. Hence $u \approx_{\lceil \mathcal{L} \rceil} v$.

For a trace $u \in \mathcal{L}_\perp$ it is immediate that $u\pi \notin \lceil \mathcal{L} \rceil$, for any π . Therefore also $u \in \lceil \mathcal{L} \rceil_\perp$. Conversely, if $v \in \lceil \mathcal{L} \rceil_\perp$, then for every path π it follows that $v\pi \notin \lceil \mathcal{L} \rceil$ and, since $\mathcal{L} \subseteq \lceil \mathcal{L} \rceil$, also $v\pi \notin \mathcal{L}$. Thus $v \in \mathcal{L}_\perp$. \blacktriangle

Intuitively, this proposition states that every property can be defined by first choosing the set of finite irremediable traces (safety closure), and then refining this safety property by excluding some infinite paths (intersection with liveness). The following characterization of liveness follows immediately from the definition of \mathcal{L}_\perp .

Lemma 28. \mathcal{L} is a liveness property if and only if $\mathcal{L}_\perp = \emptyset$. \blacktriangle

5.2 Decomposition in LTL

For an LTL formula φ we define the sets

$$C_i = \{t \in \Sigma^+ \mid t \models P_i\},$$

where P_i is the past part of the i -th conjunct of a canonical separation of anchored φ . For an index i and a trace $t \in C_i$, we

have $t\pi, 0 \models \varphi$ iff $\pi, 0 \models F_i$. Thus, if $t' \in C_i$, it follows that $t \approx_{L(\varphi)} t'$. Since a canonical separation cannot have two conjuncts with equivalent future parts, if $t \in C_i$ and $t' \notin C_i$ it follows that $t \not\approx_{L(\varphi)} t'$. This proves the following theorem; recall that \perp is the distinguished index that refers to the conjunct where $F_\perp \equiv \perp$.

Theorem 29. Let φ be an LTL formula and let $L(\varphi)$ be the property expressed by φ . The equivalence classes of the relation $\approx_{L(\varphi)}$ are given by the sets C_i and $\approx_{L(\varphi)}$ is of finite index. Furthermore, $L(\varphi)_\perp = C_\perp$. \blacktriangle

Example 30. The precedence property expressed by $\varphi = \Box(a \rightarrow \blacklozenge b)$ can be written as the following ω -regular expression.

$$L(\varphi) = (\neg a \wedge \neg b)^* b \Sigma^\omega + (\neg a \wedge \neg b)^\omega.$$

Consider the sets $C_1 = (\neg a \wedge \neg b)^* b \Sigma^*$, $C_2 = (\neg a \wedge \neg b)^+$, and $C_3 = (\neg a \wedge \neg b)^*(a \wedge \neg b) \Sigma^*$. It is straightforward to prove that these sets correspond to the equivalence classes of the relation $\approx_{L(\varphi)}$. Furthermore, $C_3 = L(\varphi)_\perp$, because every trace in C_3 is clearly irremediable. In particular, from Lemma 28 it follows that φ is not a liveness formula. Note that the equivalence classes correspond to the past parts of the three formulas of the canonical separation of Example 15. Δ

Theorem 29 establishes a connection between the relation $\approx_{L(\varphi)}$ and any canonical separation of anchored φ . Constructions of the safety closure and the safety-liveness decomposition theorem for LTL formulas are now immediate consequences of Lemma 26 and Theorem 29.

Theorem 31. Let φ be an LTL formula. The safety closure of the property $L(\varphi)$ is the property $L(\lceil \varphi \rceil)$, expressed by the LTL formula $\lceil \varphi \rceil := \Box \neg P_\perp$.

Proof. Take the correspondence of Lemma 26. We claim that $\{\pi \in \Sigma^\omega \mid \pi, 0 \models \Box \neg P_\perp\} = \{\pi \in \Sigma^\omega \mid \forall i \in \mathbb{N}. \pi_i \notin L(\varphi)_\perp\}$. The proof is straightforward: $\pi, 0 \models \Box \neg P_\perp$ iff $\forall i \in \mathbb{N}. \pi_i \models \neg P_\perp$ iff $\forall i \in \mathbb{N}. \pi_i \not\models P_\perp$ iff $\forall i \in \mathbb{N}. \pi_i \notin C_\perp$. From Theorem 29, this is equivalent to $\forall i \in \mathbb{N}. \pi_i \notin L(\varphi)_\perp$. \blacktriangle

The following lemma establishes that if a trace is irremediable for some formula, it cannot also be irremediable for its negation.

Lemma 32. For any LTL formula φ , $L(\neg\varphi)_\perp \subseteq \overline{L(\varphi)_\perp}$.

Proof. Let $t \in L(\neg\varphi)_\perp$. Then, for an arbitrary path π , $t\pi, 0 \not\models \neg\varphi$ which is equivalent to $t\pi, 0 \models \varphi$. It is immediate that $t \notin L(\varphi)_\perp$. \blacktriangle

Note that the converse inclusion is generally false: take $\varphi = \Box a$ and $t = aa$. For the path $\pi_1 = a^\omega$, we have $t\pi_1, 0 \models \varphi$ and therefore $t \in \overline{L(\varphi)_\perp}$. In contrast, for the path $\pi_2 = b^\omega$, we have $t\pi_2, 0 \models \neg\varphi$ and therefore $t \notin L(\neg\varphi)_\perp$.


Theorem 33. Let φ be an LTL formula. The formula $[\varphi] := \varphi \vee \neg[\varphi]$ is a liveness formula.

Proof. It is immediate that $L([\varphi])_\perp = \overline{L(\varphi)_\perp} \cap L(\neg[\varphi])_\perp$. From Lemma 32, it follows that $L(\neg[\varphi])_\perp \subseteq \overline{L([\varphi])_\perp}$. Therefore, from Proposition 27, it follows that $L(\neg[\varphi])_\perp \subseteq \overline{L([\varphi])_\perp}$. Thus, $L([\varphi])_\perp = \emptyset$ and by Lemma 28, the formula $[\varphi]$ is liveness. \blacktriangle

We have used canonical separation to construct the formulas $\lceil \varphi \rceil$ and $[\varphi]$, for any LTL formula φ . We have shown that $\lceil \varphi \rceil$ is a safety formula and that $[\varphi]$ is liveness. It is immediate that $\lceil \varphi \rceil \wedge [\varphi] \equiv \varphi$. This establishes the following.

Corollary 34 (Safety-Liveness Decomposition in LTL). Every LTL formula φ is initially equivalent to the conjunction of a safety LTL formula and a liveness LTL formula, given by $\lceil \varphi \rceil$ and $[\varphi]$, respectively. \blacktriangle

Figure 2. Deterministic Büchi automaton recognizing $\Box(a \vee (b \mathcal{S} c))$



These results can be used to construct a Büchi automaton accepting $L(\lceil\varphi\rceil)$. We first use the Myhill-Nerode theorem [N58] to construct a minimal DFA recognizing all the equivalence classes except C_\perp . Note that by changing the acceptance condition of an automaton $\mathcal{A} = (Q, q_0, \Delta, F)$ on finite words, we can see it as a Büchi automaton. To satisfy the conditions of the Myhill-Nerode theorem, we extend $\approx_{L(\varphi)}$ to encompass the empty word ϵ as follows. If there is an index i such that $\varphi \equiv F_i$, then we define $\epsilon \in C_i$; otherwise we define $\epsilon \notin_{L(\varphi)} t$ for every nonempty trace $t \in \Sigma^+$.

Proposition 35. *Let φ be an LTL formula and \mathcal{M} be the minimal DFA recognizing the union of equivalence classes of $\approx_{L(\varphi)}$ other than C_\perp . Then the automaton \mathcal{M} considered as a Büchi automaton recognizes $L(\lceil\varphi\rceil)$.* ▲

Corollary 36. *Let φ be a safety LTL formula and \mathcal{M} be the minimal DFA recognizing the union of equivalence classes of $\approx_{L(\varphi)}$ other than C_\perp . The automaton \mathcal{M} considered as a Büchi automaton recognizes $L(\varphi)$. Furthermore, it is minimal in the sense that it is the deterministic Büchi automaton recognizing $L(\varphi)$ with the least number of states.*

Proof. Since φ is safety, it follows that $\lceil\varphi\rceil \equiv \varphi$ and \mathcal{M} recognizes $L(\varphi)$, by Proposition 35.

For any deterministic Büchi automaton recognizing φ and any trace $t \in \Sigma^+$, the set Q_t is a singleton set. From Corollary 18, such an automaton must have at least n states, since $(P \rightarrow \bigcirc F)_n$ has n conjuncts. Furthermore, if $\varphi \not\equiv F_i$ for every index i , then there is no trace t such that $Q_t = \{q_0\}$, and the automaton must have at least $n+1$ states. This is exactly the number of states of \mathcal{M} , proving that it is minimal. ▲

Example 37. *Consider the formula $\varphi = \Box(a \vee (b \mathcal{S} c))$. For example, the property expressed by φ could state that access to a resource is not allowed (a) unless an authorized user logged into the system (c) and has been active since (b). Without going into the details of the separation procedure, we separate anchored φ , and use Construction 14 to obtain a canonical separation. The canonical separation of anchored φ consists of three conjuncts. The past and future parts are given below.*

- $P_1 = \blacksquare(a \vee (b \mathcal{S} c)) \wedge \neg(b \mathcal{S} c)$,
- $F_1 = (\Box a \vee (a \mathcal{U} c)) \wedge \bigcirc \Box(b \vee \Box a \vee (a \mathcal{U} c))$,
- $P_2 = \blacksquare(a \vee (b \mathcal{S} c)) \wedge (b \mathcal{S} c)$,

- $F_2 = \Box(b \vee \Box a \vee (a \mathcal{U} c))$,
- $P_\perp = \neg \blacksquare(a \vee (b \mathcal{S} c))$,
- $F_\perp = \perp$.

The safety closure of φ is given by $\Box \neg P_\perp \equiv \Box \blacksquare(a \vee (b \mathcal{S} c))$. It is easy to verify that this formula is initially equivalent to φ , thus showing that φ is a safety formula. It is also simple to show that $F_1 \equiv \varphi$. Therefore, the equivalence classes of the relation $\approx_{L(\varphi)}$ are given by the following regular expressions.

- $C_1 = ((a + cb^*)^*(a \wedge \neg b \wedge \neg c) + \epsilon)(a \wedge \neg c)^*$,
- $C_2 = (a + cb^*)^*cb^*$,
- $C_\perp = C_1 + C_2$.

The construction of the minimal deterministic Büchi automaton recognizing φ is now simple. The set of states is $Q = \{p_1, p_2, p_\perp\}$, where the index determines the corresponding equivalence class. The set of accepting states is $F = \{p_1, p_2\}$, and since $\epsilon \in C_1$, the initial state is p_1 . Now, for each $a \in \Sigma$ and every state q_i , we take a trace $t \in C_i$ and define $\Delta(p_i, a) = p_j$ iff $ta \in C_j$. The automaton is given in Fig. 2. ▲

6. Characterizing Classes of Temporal Properties

In this section, we use canonical separation of anchored formulas to characterize liveness, safety, stable, absolute liveness, and fairness properties. We also show that recognizing properties in these classes in LTL can be reduced to deciding LTL validity.

We start with characterizing safety and liveness in LTL.

Theorem 38 (Characterization of Liveness). *A formula φ is liveness if and only if the formula $\Diamond P_\perp$ is not satisfiable, i.e. $\Diamond P_\perp \equiv \perp$.*

Proof. From Lemma 28 and Theorem 29 it follows that φ is liveness iff every trace $t \in \Sigma^+$ falsifies P_\perp . This is equivalent to the condition that, for every path $\pi \in \Sigma^\omega$ and every $i \in \mathbb{N}$, $\pi, i \models \neg P_\perp$. Thus, φ is liveness iff $\pi, 0 \models \Box \neg P_\perp$ for every path π , and therefore also $\pi, 0 \not\models \Diamond P_\perp$. ▲

As a consequence, we obtain for LTL the following well-known result from [AS85].

Corollary 39. *A formula φ is liveness if and only if $\lceil\varphi\rceil \equiv \perp$.* ▲

A formula φ is safety if and only if the property $L(\varphi)$ is equal to $L(\Box \neg P_\perp)$. In other words, φ is safety iff $\varphi \equiv \lceil\varphi\rceil$. From Theorem 31 it follows that $L(\varphi) \subseteq L(\Box \neg P_\perp)$. Therefore, it suffices to show that there is no path $\pi \in \Sigma^\omega$ such that $\pi, 0 \models \Box \neg P_\perp$ and $\pi, 0 \not\models \varphi$. This gives us the following characterization of safety.

Theorem 40 (Characterization of Safety). *A formula φ is safety if and only if the formula $\neg\varphi \wedge \Box \neg P_\perp$ is not satisfiable.* ▲

To illustrate these results, we return to our running example.

Example 41. *Consider again $\varphi = \Box(a \rightarrow \blacklozenge b)$. In Example 30, we gave a canonical separation of anchored φ . It is immediate that $P_\perp = \blacklozenge(a \wedge \blacksquare \neg b)$. Consider the formula $\blacklozenge(a \wedge \blacksquare \neg b)$. This formula is satisfiable, for example by the path a^ω . Therefore, φ is not liveness. We can explicitly construct the safety closure of φ as*

$$\lceil\varphi\rceil = \Box \neg \blacklozenge(a \wedge \blacksquare \neg b).$$

This formula is globally equivalent to $\Box \blacksquare(a \rightarrow \blacklozenge b)$, and initially equivalent to $\Box(a \rightarrow \blacklozenge b)$. Therefore, the formula $\neg\varphi \wedge \lceil\varphi\rceil$ is not satisfiable, proving that φ is a safety property. ▲

We next characterize stable, absolute liveness, and fairness properties in LTL. We follow Sistla's definitions for these properties [S94].

Definition 42. Let φ be an LTL formula.

1. φ is **stable** if for every path π , the initial satisfaction $\pi, 0 \models \varphi$ implies $\pi^i, 0 \models \varphi$, for all $i \in \mathbb{N}$.
2. φ is **absolute liveness** if φ is satisfiable, and for all traces t and every path π such that $\pi, 0 \models \varphi$, then $t\pi, 0 \models \varphi$.
3. φ is **fairness** if it is both stable and absolute liveness. \triangle

Intuitively, a fairness property can neither be satisfied nor falsified in finite time. Note that Sistla's definition of stable properties, stating that $\pi, 0 \models \varphi$ implies $\pi, i \models \varphi$, differs slightly from ours. The two definitions coincide for FLTL, but the modification above is necessary to formalize a stable property as a property containing all suffixes of its elements in LTL.

Let φ be an LTL formula and let $(P \rightarrow \bigcirc F)_n$ be a canonical separation of anchored φ . We first give necessary and sufficient conditions for φ to be stable and absolute liveness. We combine them to characterize fairness and conclude the section with an example.

For any past formula ψ , we can find an initially equivalent propositional formula. We just iteratively replace $\psi_1 \mathcal{S} \psi_2$ by ψ_2 and $\blacklozenge \psi$ by \perp . For every past part P_i , let us denote this propositional counterpart by N_i .

The intuition for the following theorems is as follows. Let π be a path satisfying a stable formula φ and consider a canonical separation of anchored φ . Then for any k there is an index i such that $\pi, k \models P_i$ and $\pi, k \models \bigcirc F_i$. If we remove a prefix of π of length k , it is possible that $\pi^k, 0 \models P_j$, i.e. $\pi, k \models N_j$, for some $j \neq i$. In this case, it is necessary that $\pi, k \models \bigcirc F_j$ for φ to be stable. Analogous intuition holds for absolute liveness

Theorem 43. A formula φ is stable if and only if the formula $\Psi_1(\varphi) = \bigvee_{i,j=1}^n \diamond(P_i \wedge N_j \wedge \bigcirc F_i \wedge \neg \bigcirc F_j)$ is not satisfiable.

Proof. Let φ be stable and assume the contrary, i.e. there is a path $\pi \in \Sigma^\omega$ such that $\pi, 0 \models \Psi_1(\varphi)$. Then there exists a $k \geq 0$ and indices i, j such that $\pi, k \models P_i$, $\pi, k \models N_j$, $\pi, k \models \bigcirc F_i$, and $\pi, k \models \neg \bigcirc F_j$. Since $\pi, k \models P_i$ and $\pi, k \models \bigcirc F_i$, it follows that $\pi, 0 \models \varphi$. From $\pi, k \models N_j$, it follows that $\pi^k \models P_j$. Since $\pi, k \models \neg \bigcirc F_j$, we also have $\pi^k, 0 \models \neg \bigcirc F_j$, and $\pi^k, 0 \not\models \varphi$. This contradicts the assumption that φ is stable. \blacktriangle

To prove the converse, suppose $\Psi_1(\varphi)$ is not satisfiable and φ is not stable. Then, there exists a path π such that $\pi, 0 \models \varphi$ and a $k > 0$ such that $\pi^k, 0 \not\models \varphi$. Let i and j be indices such that $\pi_k \models P_i$ and $\pi_k \models N_j$. It follows that $\pi^k, 0 \models \bigcirc F_i$ and $\pi^k, 0 \not\models \bigcirc F_j$. Therefore, $\pi, k \models P_i \wedge N_j \wedge \bigcirc F_i \wedge \neg \bigcirc F_j$, and $\pi, 0 \models \Psi_1(\varphi)$, which is a contradiction. \blacktriangle

Example 44. Let $\varphi = \Box a$. A canonical separation of anchored φ is $(\blacklozenge \neg a \rightarrow \bigcirc \perp) \wedge (\blacksquare a \rightarrow \bigcirc \Box a)$. To prove φ is a stable formula, by Theorem 43 we show that the formulas $\diamond(\blacklozenge \neg a \wedge a \wedge \bigcirc \perp \wedge \neg \bigcirc \Box a)$ and $\diamond(\blacksquare a \wedge \neg a \wedge \bigcirc \Box a \wedge \bigcirc \top)$ are not satisfiable. This is immediate because of the $\bigcirc \perp$ in the first formula and $\blacksquare a \wedge \neg a$ in the second formula. \triangle

The proof of the following theorem is similar to the proof of Theorem 43.

Theorem 45. A formula φ is absolute liveness if and only if the formula $\Psi_2(\varphi) = \bigvee_{i,j=1}^n \diamond(P_i \wedge N_j \wedge \neg \bigcirc F_i \wedge \bigcirc F_j)$ is not satisfiable and φ is satisfiable. \blacktriangle

Example 46. Let $\varphi = \diamond a$. A canonical separation of anchored φ is $(\perp \rightarrow \bigcirc \perp) \wedge (\blacklozenge a \rightarrow \bigcirc \top) \wedge (\blacksquare \neg a \rightarrow \bigcirc \diamond a)$. It is immediate that φ is a liveness formula, and therefore satisfiable. To prove it is absolute liveness, by Theorem 45 we show that $\Psi_2(\varphi)$ is not satisfiable. Since $P_\perp = \perp$, we need only consider indices $i, j \neq \perp$. The two formulas we must check are $\diamond(\blacklozenge a \wedge \neg a \wedge \bigcirc \perp \wedge \bigcirc \diamond a)$ and

Table 1. Reduction to UNSAT. A canonical separation of anchored φ is given by $(P \rightarrow \bigcirc F)_n$. Note*: for absolute liveness, φ should also be liveness.

Property φ	Characterization (UNSAT)
Safety	$\neg \varphi \wedge \Box \neg P_\perp$
Liveness	$\diamond P_\perp$
Absolute Liveness *	$\bigvee_{i,j=1}^n \diamond(P_i \wedge P_j^F \wedge \neg F_i \wedge F_j)$
Stable	$\bigvee_{i,j=1}^n \diamond(P_i \wedge P_j^F \wedge F_i \wedge \neg F_j)$
Fairness	$\diamond \neg P_i$, for some i

$\diamond(\blacksquare \neg a \wedge a \wedge \neg \bigcirc \diamond a \wedge \bigcirc \top)$. It is immediate that neither formula is satisfiable. \triangle

Combining the characterizations of stable and absolute liveness properties results in a simple characterization for fairness.

Theorem 47. A formula φ is fairness if and only if a canonical separation of anchored φ is of the form $(\perp \rightarrow \bigcirc \perp) \wedge (\top \rightarrow \bigcirc F)$, for some future formula F .

Proof. Let φ be fairness. In particular, φ is liveness, satisfiable, and $\diamond P_\perp \equiv \perp$. Assume that there are indices $i \neq j$ and traces $u, v \in \Sigma^+$ such that $u \models P_i$ and $v \models P_j$. The formula F_i is then satisfiable and let π be a path such that $\pi, 0 \models F_i$. Then $u\pi, 0 \models \varphi$. Since φ is stable, it follows that $\pi, 0 \models \varphi$. Since φ is absolute liveness, $v\pi, 0 \models \varphi$. Therefore $\pi, 0 \models F_j$. By repeating this proof with the indices exchanged, we get that F_i and F_j are equivalent, contradicting the initial assumption. \blacktriangle

For the converse, the canonical separation satisfies the conditions of Theorem 43 and Theorem 45. It is thus immediate that φ is both stable and absolute liveness. \blacktriangle

Corollary 48 (Characterization of Fairness). A formula φ is fairness iff it is satisfiable and there is an index i such that $\diamond \neg P_i \equiv \perp$. \blacktriangle

Corollary 49. A formula φ is fairness if and only if a canonical separation of anchored φ is of the form $(\perp \rightarrow \bigcirc \perp) \wedge (\top \rightarrow \bigcirc F)$, where $F \equiv \varphi$. \blacktriangle

This characterization implies that the LTL formula of Example 20 expresses a fairness property.

Corollary 50. Recognizing safety, liveness, absolute liveness, stable and fairness properties can be reduced to the UNSAT problem for LTL, summarized in Table 1. \blacktriangle

A formula φ of LTL expresses a safety, liveness, absolute liveness, stable or fairness property iff the respective formula in Table 1 is not satisfiable. Note that the characterization of absolute liveness additionally requires that φ is satisfiable, but one can instead show that φ is liveness. This is justified since every absolute liveness formula is a liveness formula, and every liveness formula is satisfiable.

7. Related Work

The canonical separation of anchored formulas is inspired by the (declarative) past implies (imperative) future paradigm of [G87], which was used for the separated normal form (SNF) of Fisher [F97] and is the underlying idea of MetateM [F92]. Intuitively, SNF is a set of rules that describe which actions a system has to make at the current time with respect to observations it has made in the previous steps. Our reasoning is based on finite traces, while with SNF one can only reason about the current moment in time, which makes it necessary to start at the initial time 0.

Concerning the complexity of translating LTL to FLTL, a nonelementary algorithm can be obtained using the separation rewrite rules from [G87]. Hodkinson and Reynolds [HR05] state that it is not clear whether this can be done without separation. An elementary, but inefficient, algorithm for the translation certainly exists if the succinctness gap is elementary. One can just check initial equivalence with respect to FLTL formulas of increasing size until a match is found.

It appears that the existence of an elementary bound on the succinctness gap has become folklore. Since ω -automata recognizing the language of any LTL formula can be efficiently constructed, an elementary translation from ω -automata to FLTL would answer the question. Several authors have proposed using a translation by Wilke [W99] for this purpose. However, Wilke considers FLTL on finite traces and generalizing the translation to FLTL on infinite paths is not obvious. Gastin and Oddoux [GO01] generalize Wilke's result to infinite words, but their construction uses first order logic as an intermediary step, which incurs a nonelementary blow-up. Maler and Pnueli [MP94] give a translation from Muller automata to LTL, but they use past temporal connectives. These formulas can be translated to FLTL using Pnueli and Zuck's algorithm [PZ93]; the algorithm is however nonelementary.

Sistla [S94, S85] characterizes safety, stable, absolute liveness, fairness, and other properties in FLTL, leaving the characterization of liveness as an open problem. He also gives a decision procedure for recognizing safety and liveness in FLTL and characterizes various properties in less expressive sub-logics of FLTL. Lichtenstein et al. [LPZ85] characterize safety in LTL, but their characterizations do not result in a decision procedure. They also characterize "liveness", but their definition of liveness differs from Alpern and Schneider's (which we follow here). Manna and Pnueli [MP87] use Gabbay's separation to characterize liveness properties in LTL, resulting in a characterization similar to ours. We are not aware of any characterizations of stable, absolute liveness, and fairness properties in LTL.

Concerning the safety-liveness decomposition, Alpern and Schneider [AS85] prove that every property is an intersection of a safety and a liveness property. Moreover, they show that ω -regular properties, i.e. the properties accepted by Büchi automata, can be decomposed as an intersection of a safety ω -regular property and a liveness ω -regular property [AS87]. Their decomposition does not readily translate to LTL. This is because LTL formulas are only capable of expressing star-free ω -regular languages, which constitute a strict subset of ω -regular languages [W83].

In previous work [PMTDB], we prove that the safety-liveness decomposition is possible in FLTL using counter-free Büchi automata. The argument relies on a special case of a translation from counter-free DFA to FLTL on finite words [W99] and the notion of protected formulas from [PZ93]. The construction of the safety closure we present here is elementary, in contrast to the nonelementary construction in [PMTDB].

Acknowledgments

This work was partially supported by the FP7-ICT-2009-5 Project no. 257876, "Secure Provision and Consumption in the Internet of Services" (<http://www.spacios.eu>). The authors would like to thank Eugen Zălinescu and Ognjen Marić for valuable feedback and suggestions.

References

[AS85] B. Alpern, and F. B. Schneider. Defining Liveness, *Information Processing Letters* 21(4):181-185, 1985.

- [AS87] B. Alpern, and F. B. Schneider. Recognizing Safety and Liveness, *Distributed Computing* 2(3):117-126, 1987.
- [F97] M. Fisher. A Normal Form for Temporal Logics and its Applications in Theorem-Proving and Execution, *J. Log. Comput.* 7(4):429-456, 1997.
- [F92] M. Fisher, and P. Noël. Transformation and Synthesis in MetateM Part I: Propositional MetateM, Technical report, Department of Computer Science, University of Manchester, 1992.
- [G87] D. M. Gabbay. The Declarative Past and Imperative Future: Executable Temporal Logic for Interactive Systems, *Temporal Logic in Specification, Lecture Notes in Computer Science* 398:409-448. Springer, 1987.
- [GPSS80] D. M. Gabbay, A. Pnueli, S. Shelah, and J. Stavi. On the Temporal Basis of Fairness, *POPL*:163-173. ACM Press, 1980.
- [GO01] P. Gastin, and D. Oddoux. Fast LTL to Büchi Automata Translation, *CAV, Lecture Notes in Computer Science* 2102:53-65. Springer 2001.
- [HR05] I. M. Hodkinson, and M. Reynolds, Separation - Past, Present, and Future, *We Will Show Them!* (2):117-142. College Publications, 2005.
- [H79] J. E. Hopcroft, and J. D. Ullman, *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley, 1979.
- [L77] L. Lamport, Proving the Correctness of Multiprocess Programs, *IEEE Trans. Software Eng.* 3(2):125-143, 1977.
- [LPZ85] O. Lichtenstein, A. Pnueli, and L. D. Zuck. The Glory of the Past, *Logic of Programs, Lecture Notes in Computer Science* 193:196-218. Springer, 1985.
- [MP90] O. Maler, and A. Pnueli. Tight Bounds on the Complexity of Cascaded Decomposition of Automata, *FOCS*:672-682. IEEE Computer Society, 1990.
- [MP94] O. Maler, and A. Pnueli. On the Cascaded Decomposition of Automata, its Complexity and its Application to Logic, *ACTS Mobile Communication* 48, 1994.
- [MP87] Z. Manna, and A. Pnueli. A Hierarchy of Temporal Properties, In *Proceedings of the Sixth Annual ACM Symposium on Principles of Distributed Computing, PODC '87*. ACM, 1987.
- [MP89] Z. Manna, and A. Pnueli. The anchored version of the temporal framework, *Linear Time, Branching Time and Partial Order in Logics and Models for Concurrency, LNCS* 354:201-284. Springer, 1989.
- [M03] N. Markey. Temporal logic with past is exponentially more succinct, *Concurrency Column, Bulletin of the EATCS* 79:122-128, 2003.
- [M71] R. McNaughton, and A. Papert. *Counter-Free Automata*, M.I.T. research monograph no. 65. The MIT Press, 1971.
- [M74] A. Meyer. The inherent complexity of theories of ordered sets, In *the Proceedings of the International Congress of Mathematics*:477-482, 1974.
- [N58] A. Nerode. Linear Automaton Transformations, In *AMS* 9. AMS, 1958.
- [PMTDB] G. Petric Maretić, M. Torabi Dashti, and D. Basin. LTL is Closed Under Topological Closure, *Information Processing Letters*, 114(8):408-413. Elsevier, 2014.
- [PZ93] A. Pnueli, and L. D. Zuck. In and Out of Temporal Logic, *LICS*:124-135. IEEE Computer Society, 1993.
- [S85] A. P. Sistla. On Characterization of Safety and Liveness Properties in Temporal Logic, *PODC*:39-48. ACM, 1985.
- [S94] A. P. Sistla. Safety, Liveness and Fairness in Temporal Logic, *Formal Asp. Comput.* 6(5):495-512, 1994.
- [T90] W. Thomas. Automata on Infinite Objects, *Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics*:133-192. Elsevier and MIT Press, 1990.
- [W99] T. Wilke. Classifying Discrete Temporal Properties, *STACS* 99, *Lecture Notes in Computer Science* 1563:32-46. Springer, 1999.
- [W83] P. Wolper. Temporal Logic Can Be More Expressive, *Information and Control* 56(1/2):72-99, 1983.