

Size-Change Termination and Satisfiability for Linear-Time Temporal Logics

Martin Lange

School of Electr. Eng. and Computer Science, University of Kassel, Germany

Abstract. In the automata-theoretic framework, finite-state automata are used as a machine model to capture the operational content of temporal logics. Decision problems like satisfiability, subsumption, equivalence, etc. then translate into questions on automata like emptiness, inclusion, language equivalence, etc. Linear-time temporal logics like LTL, PSL and the linear-time μ -calculus have relatively simple translations into alternating parity automata, and this automaton model is closed under all Boolean operations with very simple constructions. Thus, the typical decision problems for such linear-time temporal logics reduce relatively simply to the emptiness problem for alternating parity automata. In this paper we present a method for decision this emptiness problem without going through intermediate automaton models like nondeterministic ones. The method is a direct adaptation of the size-change termination principle which was originally used to decide termination of abstract functional programs.

1 Introduction

Temporal logics are some of the most well-known and established tools for the specification of systems evolving in time. In computer science, they are mainly interesting as formal languages used to describe, reason about, analyse and verify program behaviour [7].

Temporal logics come in two different kinds depending on the nature of time underlying the models that they get interpreted about: linear-time and branching-time [20, 21, 25]. The viewpoint of linear time is that every moment in time has a unique successor, i.e. the future is determined. In branching time, a moment may have several successors. Here we only deal with linear time, namely we consider the well-known simple linear-time temporal logic LTL [16], as well as two of its extensions: a core of the industry standard property specification language PSL [2] which extends LTL with semi-regular expressions in order to remedy LTL's weakness of not being able to define all ω -regular properties; as well as the linear-time μ -calculus [22, 3] which uses second-order quantification in the form of least and greatest fixpoints in order to obtain higher expressivity.

A prominent methodology for obtaining decision procedures is the automata-theoretic framework. It is particularly suitable for linear-time logics since their models can be seen as infinite words which immediately links logics and automata as two different specification formalisms for languages of such words. Logics are

often more natural to use as specification languages but automata are often closer to a decision procedure. Hence, translations from formulas to automata preserving their set of models are desirable, and they exist for the aforementioned logics.

It is known that every LTL formula can be translated into a nondeterministic Büchi automaton [26] with an exponential blow-up. The same holds for PSL [4] although the blow-up is in general doubly exponential. LT_μ can also be translated into such automata at a singly exponential blow-up only [26].

Translations into nondeterministic automata are particularly useful in order to decide satisfiability problems because satisfiability on the logical side corresponds to non-emptiness on the automata side, and non-emptiness problems for nondeterministic automata are usually solved via simple reachability problems on graphs. Other problems, however, in particular the universality and inclusion problem are as difficult for nondeterministic automata as the satisfiability problem for the corresponding temporal logics is. Note that on the logical side, problems like validity, subsumption and equivalence easily reduce to the satisfiability problem. This has led to the use of a richer automaton model: alternating automata. They typically enable a simple translation from formulas and on top of that a more difficult decision procedure as opposed to nondeterministic ones which usually come with a difficult translation and then simpler decision procedures. If “simple” means “polynomial” and “difficult” means “exponential” then the route via alternating automata may even be better in terms of complexity.

Translations from the linear-time temporal logics mentioned above into alternating automata are known [23, 24, 14, 10]. In order to obtain decision procedures for these logics, one then only needs decision procedures for the corresponding problems on alternating automata. In fact, it suffices to be able to solve the emptiness problem for alternating automata just like it suffices to solve the satisfiability problem for temporal logics in order to solve all sorts of other problems through simple reductions.

The standard way to solve the emptiness problem for alternating automata has been using translations into nondeterministic automata. It may be the surprising simplicity of the Miyano-Hayashi construction [15] translating alternating Büchi into nondeterministic Büchi automata in comparison to the problem of turning a nondeterministic one into a deterministic one, that has put a brake onto research on different and possibly direct methods for the emptiness problem for alternating automata. This construction can be generalised to richer acceptance conditions like Streett automata [6], yet it still aims at translating into a nondeterministic model first.

Here we propose a different and direct method for the analysis of the emptiness problem for alternating parity automata. It originates from termination analysis for abstract functional programs and is called *size-change termination* (SCT) [13]. It is noted that the problem underlying this particular termination analysis can be solved by a reduction to the inclusion problem for nondeterministic Büchi automata but, since this requires complementation, SCT is proposed.

SCT has not passed unnoticed in the world of temporal logics and automata on infinite words: it has first been used to decide validity for the linear-time μ -calculus [5] and then prosed as a method to decide universality and inclusion for nondeterministic Büchi automata [8, 9, 1]. In fact, the first real work in this area is Büchi's original complementation proof for nondeterministic Büchi automata since the decision problems based on SCT use the same techniques. They are often called *Ramsey-based* because their correctness proof usually relies on the famous combinatorial Ramsey Theorem [17].

This paper is organised as follows. In Sect. 2 we recall the three important temporal logics mentioned above: LTL, PSL and the linear-time μ -calculus. In Sect. 3 we recall alternating parity automata and various subclasses thereof, and sketch how their emptiness problems characterise typical problems like satisfiability, subsumption, etc. for the temporal logics at hand. In Sect.4 we describe an SCT based method to decide emptiness of alternating parity automata.

2 Linear-Time Temporal Logics

2.1 Infinite Words

Let $\mathcal{P} = \{p, q, \dots\}$ be a countably infinite set of *atomic propositions*. Linear-time temporal logics are interpreted over ω -sequences of sets of such propositions: an *infinite word* w is an element of $(2^{\mathcal{P}})^{\omega}$. A *finite word* is a $v \in (2^{\mathcal{P}})^*$. We write ϵ for the empty word of length 0, and $|v|$ in general for the length of the finite word v .

If w is a word $A_0A_1\dots$ for $A_i \subseteq \mathcal{P}$ then $w(i)$ is used to denote A_i . We write $w(i, j)$ to denote the finite subsequence $A_i \dots A_j$. Note that $w(i, j) = \epsilon$ if $j < i$.

2.2 LTL

One of the simplest and most well-known linear-time temporal logics is LTL. Its formulas are built from atomic propositions using Boolean operators and two temporal operators: the *next state* operator \bigcirc , and the *until* operator \bigcup . Formulas of LTL are given by the following grammar.

$$\varphi ::= q \mid \varphi \wedge \psi \mid \neg\varphi \mid \bigcirc\varphi \mid \varphi \bigcup \psi$$

Formulas of LTL are interpreted in positions i of an infinite word as follows.

$$\begin{aligned} w, i \models q & \quad \text{iff } q \in w(i) \\ w, i \models \varphi \wedge \psi & \quad \text{iff } w, i \models \varphi \text{ and } w, i \models \psi \\ w, i \models \neg\varphi & \quad \text{iff } w, i \not\models \varphi \\ w, i \models \bigcirc\varphi & \quad \text{iff } w, i+1 \models \varphi \\ w, i \models \varphi \bigcup \psi & \quad \text{iff there is } j \geq i \text{ s.t. } w, j \models \psi \text{ and for all } h \text{ with } i \leq h < j \\ & \quad \text{we have } w, h \models \varphi \end{aligned}$$

Further Boolean operators like disjunction $\varphi \vee \psi$, implication $\varphi \rightarrow \psi$, are defined as abbreviations in the usual way. Other temporal operators can also be derived, for instance $\mathbf{F}\varphi := \mathbf{true} \mathbf{U} \varphi$ (“finally”) where $\mathbf{true} := q \vee \neg q$ for some $q \in \mathcal{P}$; $\mathbf{G}\varphi := \neg \mathbf{F} \neg \varphi$ (“generally”), etc.

We write $|\varphi|$ for the size of the formula φ , measured in terms of its DAG representation or, equivalently, the number of different subformulas.

Example 1. LTL can easily express properties concerning infinite occurrences of some atomic proposition in a word. For example, $\mathbf{GF}q \wedge \mathbf{FG}\neg p$ expresses that q holds in infinitely many positions and p holds almost everywhere, i.e. in all but finitely many positions.

A *model* for φ is a $w \in (2^{\mathcal{P}})^{\omega}$ s.t. $w, 0 \models \varphi$. We write $L(\varphi)$ for the set of all models of φ . The *satisfiability problem* is: given a formula φ , decide whether or not there is a model for it, i.e. whether or not $L(\varphi) \neq \emptyset$.

Proposition 1 ([19]). *The satisfiability problem for LTL is PSPACE-complete.*

We also describe two extensions of LTL in the following. They differ in their syntax and their semantics need more technicalities, but central concepts like formula size as well as set of models are defined as they are for LTL. Hence, we will not repeat them explicitly anymore.

2.3 PSL – An Extension of LTL

It is known that LTL cannot express counting properties like “ q holds in every second position of a word” [27]. Note that this is an ω -regular property. There are several ways to overcome this weakness, for instance by introducing quantification over positions, i.e. by extending LTL with stronger logical connectives. PSL extends LTL with tools from the domain of formal languages, namely semi-extended regular expressions.

The language of all *Boolean expressions* over atomic propositions as above is given by the following grammar.

$$\zeta ::= q \mid \zeta \wedge \zeta \mid \neg \zeta$$

Other Boolean operators can be defined as abbreviations. The satisfaction relation between a set $A \subseteq \mathcal{P}$ and a Boolean expression ζ is defined straightforwardly, i.e. $A \models \zeta$ iff ζ evaluates to 1 under the usual rules for Boolean connectives when all atomic propositions in A are set to 1 and all in $\mathcal{P} \setminus A$ are set to 0.

Semi-extended regular expressions (SERE) over Boolean expressions are built according to the following grammar.

$$\alpha ::= \zeta \mid \alpha \cup \alpha \mid \alpha \cap \alpha \mid \alpha; \alpha \mid \alpha^*$$

where ζ is a Boolean expression as above. We write α^+ to abbreviate $\alpha; \alpha^*$.

A SERE is interpreted in a finite subword of an infinite $w \in (2^{\mathcal{P}})^\omega$ as follows. Note that such a subword is uniquely identified by two positions $i, j \in \mathbb{N}$ with $i \leq j$.

$$\begin{aligned}
w, i, j \models \zeta & \quad \text{iff } i = j \text{ and } w(i) \models \zeta \\
w, i, j \models \alpha \cup \beta & \quad \text{iff } w, i, j \models \alpha \text{ or } w, i, j \models \beta \\
w, i, j \models \alpha \cap \beta & \quad \text{iff } w, i, j \models \alpha \text{ and } w, i, j \models \beta \\
w, i, j \models \alpha; \beta & \quad \text{iff there is } h \text{ s.t. } i \leq h \leq j \text{ and } w, i, h \models \alpha \text{ and } w, h, j \models \beta \\
w, i, j \models \alpha^* & \quad \text{iff there are } n \geq 0 \text{ and } h_0, \dots, h_n \text{ s.t. } h_0 = i, h_n = j \text{ and} \\
& \quad w, h_k, h_{k+1} \models \alpha \text{ for all } k = 0, \dots, n-1
\end{aligned}$$

Note that in the first clause, satisfaction of a SERE in a finite word is reduced to satisfaction of a Boolean expression in a symbol of that word.

Formulas of PSL are then built by extending the syntax of LTL with operators that make use of SERE. Note that the standard of PSL [2] describes many operators for the logic; here we concentrate on two of them only – the “and then” and the “closure” operator. The constructions to follow can easily be extended to cover other PSL operators as well though.

$$\varphi ::= q \mid \varphi \wedge \varphi \mid \neg \varphi \mid \varphi \mathbf{U} \varphi \mid \alpha \diamond \rightarrow \varphi \mid \mathbf{C1} \alpha$$

The interpretation in positions of a word $w \in (2^{\mathcal{P}})^\omega$ extends the one for LTL given above by two clauses.

$$\begin{aligned}
w, i \models \alpha \diamond \rightarrow \varphi & \quad \text{iff there is } j \geq i \text{ s.t. } w, i, j \models \alpha \text{ and } w, j \models \varphi \\
w, i \models \mathbf{C1} \alpha & \quad \text{iff for all } j \geq i \text{ exists } v \in (2^{\mathcal{P}})^* \text{ s.t. } w(i, j)v \models \alpha
\end{aligned}$$

Example 2. The aforementioned language $L = \{w \in (2^{\mathcal{P}})^\omega \mid \forall i \in \mathbb{N} : q \in w(2i)\}$ of all words in which q holds in every even position can be defined in PSL by the formula $\mathbf{C1}(q; \mathbf{true})^*$. Equally, it is defined by $q \wedge \neg((\mathbf{true}; \mathbf{true})^+ \diamond \rightarrow \bigcirc \neg q)$.

Proposition 2 ([4, 12]). *The satisfiability problem for PSL is EXPSPACE-complete.*

The exponential increase in complexity compared to LTL is owed to the use of the intersection operator in semi-extended regular expressions. Note that these can be translated into nondeterministic finite automata (NFA) at a blow-up that is polynomial in the size of such an SERE but exponential in the nesting depth of the intersection operator. The logic obtained by replacing SERE with NFA is closely related to LTL with automata connective which also has a PSPACE-complete satisfiability problem [11].

2.4 The Linear-Time μ -Calculus

The Linear-Time μ -Calculus LT_μ is not directly an extension of LTL. It obtains ω -regular expressive power by adding least (finite iteration) and greatest (infinite

iteration) fixpoint operators to the fragment of LTL without the until operator. Let $\mathcal{V} = \{X, Y, \dots\}$ be a countably infinite set of variable. Formulas of LT_μ are constructed as follows.

$$\varphi ::= q \mid X \mid \varphi \wedge \varphi \mid \neg\varphi \mid \bigcirc\varphi \mid \mu X.\varphi$$

where $q \in \mathcal{P}$ and $X \in \mathcal{V}$. We require that each formula is well-formed in the sense that every variable is bound by a binder μ at most once, and in every subformula $\mu X.\psi$ every occurrence of X is in the scope of an even number of negation symbols. For instance, $\mu X.\neg\mu Y.(\neg p \vee \bigcirc \neg Y) \wedge \bigcirc \neg X$ is well-formed because both variable occurrences are under two (different) nested negation operators.

Alongside the *least fixpoint quantifier* μ we introduce the *greatest fixpoint quantifier* ν via $\nu X.\varphi := \mu X.\neg\varphi[\neg X/X]$ where $\varphi[\psi/X]$ denotes the formula that is obtained from φ by replacing every free occurrence of X with ψ . Then the formula above can be written entirely without negation symbols as $\mu X.\nu Y.(p \wedge \bigcirc Y) \vee \bigcirc X$.

In order to interpret an LT_μ formula with free variables in a position of a word $w \in (2^{\mathcal{P}})^\omega$ we need the help of environments $\rho : \mathcal{V} \rightarrow 2^{\mathcal{N}}$. We write $\rho[X \mapsto P]$ for the environment that maps X to P and behaves like ρ on all other arguments.

$$\begin{aligned} w, i, \rho \models q & \quad \text{iff } q \in w(i) \\ w, i, \rho \models \varphi \wedge \psi & \quad \text{iff } w, i, \rho \models \varphi \text{ and } w, i, \rho \models \psi \\ w, i, \rho \models \neg\varphi & \quad \text{iff } w, i, \rho \not\models \varphi \\ w, i, \rho \models \mu X.\varphi & \quad \text{iff } i \in \bigcap \{P \subseteq \mathbb{N} \mid P \supseteq \{j \mid w, j, \rho[X \mapsto P] \models \varphi\}\} \end{aligned}$$

Example 3. The language of all words containing q in every even position can easily be defined in LT_μ : $\nu X.q \wedge \bigcirc \bigcirc X$.

The formula $\mu X.\nu Y.(p \wedge \bigcirc Y) \vee \bigcirc X$ mentioned above states that p holds in almost all positions. I.e. it is equivalent to the LTL formula $\text{FG}p$.

Proposition 3 ([22]). *The satisfiability problem for LT_μ is PSPACE-complete.*

The reason for introducing LT_μ is the fact that it subsumes the two important temporal logics LTL and PSL. While this is also trivially true for PSL, LT_μ provides a clean (albeit not necessarily intuitive) syntax which is advantageous for the further treatment of these logics.

Proposition 4. *There are equivalence-preserving translations from ...*

- LTL into LT_μ that incur a linear blow-up,
- PSL into LT_μ that incur a blow-up which is polynomial in the size of the formula and exponential in the nesting depth of the intersection operators [12].

The translation from LTL into LT_μ is realised by the fact that $\varphi \text{ U } \psi$ can be expressed as $\mu X.\psi \vee (\varphi \wedge \bigcirc X)$. The translation from PSL is more complicated and uses the fact that SERE can be translated into NFA, as well as a close

resemblance between LT_μ formulas and automata. This close resemblance will be used in the following where we introduce alternating parity automata, a model of finite automata operating on infinite words that easily captures LT_μ and allows several problems on other automata to be regarded as emptiness problems.

3 Automata on Infinite Words

3.1 Alternating Parity Automata

We introduce a very powerful automaton model which captures many well-known models of automata on infinite words, including (non)deterministic Büchi and co-Büchi automata.

For a set M let $\mathbb{B}^+(M)$ denote that set of all positive Boolean formulas over M , i.e. the least set that contains M and satisfies: if $\{\varphi, \psi\} \subseteq \mathbb{B}^+(M)$ then $\{\varphi \wedge \psi, \varphi \vee \psi\} \subseteq \mathbb{B}^+(M)$.

A *alternating parity automaton* (APA) is a tuple $\mathcal{A} = (Q, AP, q_0, \delta, \Omega)$ where Q is a finite set of states, AP is a finite subset of \mathcal{P} as used in the previous section, $q_0 \in Q$ is a designated starting state, $\delta : Q \rightarrow \mathbb{B}^+(Q \cup AP \cup \neg AP)$ with $\neg AP := \{\neg q \mid q \in AP\}$ is the transition function, and $\Omega : Q \rightarrow \mathbb{N}$ assigns priorities to the states.

Here we measure the size of an automaton, $|\mathcal{A}|$, as the number of its states.

A *run* of the APA \mathcal{A} on a word $w \in (2^{\mathcal{P}})^\omega$ is a leveled DAG t whose nodes are labeled with states from Q , that has a single root on level 0, and the successors of a node on level n are all on level $n + 1$. Furthermore, it obeys the following rules. We write $t(v)$ for the label of node v .

1. We have $t(v_0) = q_0$ for the root v_0 .
2. Take any node v on some level n and let u_1, \dots, u_k be the set of its successors. Then we have $\{t(u_1), \dots, t(u_k)\} \models \delta(t(v), w(n))$.

Such a run t is *accepting* if on every path through t the greatest priority of states that occur infinitely often is even. The *language* of \mathcal{A} is $L(\mathcal{A}) = \{w \mid \text{there is an accepting run of } \mathcal{A} \text{ on } w\}$.

Example 4. Take the language $L = \{w \in (2^{\{p,q\}})^\omega \mid \text{if } q \text{ holds infinitely often, then } q \wedge p \text{ holds infinitely often in } w\}$. It is accepted by the APA $(\{0, 1, 2\}, \{p, q\}, 0, \delta, \Omega)$ where Ω is the identity function, and the transition function is defined for all three states i as

$$\delta(i) = (p \wedge q \wedge 2) \vee 1 \vee (\neg q \wedge 0)$$

It is also accepted by the APA $(\{0, 1, 2\}, \{p, q\}, 0, \delta, \Omega)$ where $\Omega(0) = \Omega(2) = 0$ and $\Omega(1) = 1$, and

$$\begin{aligned} \delta(0) &= (\neg q \vee p \vee 1) \wedge 0 \\ \delta(1) &= (q \wedge p \wedge 2) \vee 1 \\ \delta(2) &= 2 \end{aligned}$$

3.2 Subclasses of Alternating Parity Automata

Let $\mathcal{A} = (Q, AP, q_0, \delta, \Omega)$ be an APA. Then \mathcal{A} is an *alternating Büchi automaton* (ABA) if $\Omega : Q \rightarrow \{1, 2\}$. It is an *alternating co-Büchi automaton* (AcoBA) if $\Omega : Q \rightarrow \{0, 1\}$. I.e. Büchi acceptance is concerned with the infinite occurrence of some states whereas co-Büchi acceptance demands that certain states occur almost everywhere in a path of a run.

The literature contains different definitions of a weak automaton, sometimes constraining the graph structure of the automaton, sometimes weakening the acceptance condition by changing “occurs infinitely often” into “occurs” simply. These notions are equivalent though [14]. Here we consider the former. \mathcal{A} is called *weak* (WAPA) if for all $q, q' \in Q$ s.t. q' occurs in $\delta(q)$ we have $\Omega(q') \leq \Omega(q)$. Hence, the priorities on every path in a run of a weak automaton are monotonically decreasing, and the largest priority that occurs infinitely often is automatically the one that occurs almost everywhere. Consequently, weak alternating automata can easily be defined as ABA or AcoBA.

\mathcal{A} is *nondeterministic* if for all $q \in Q$ we have that $\delta(q)$ is a disjunction of minterms containing exactly one state, i.e. all other conjuncts are atomic propositions or negations thereof.

3.3 Constructions on Alternating Parity Automata

Alternating automata are closed under all Boolean operations.

Proposition 5. *Let \mathcal{A} and \mathcal{B} be two APA. There are APA*

1. $\bar{\mathcal{A}}$ s.t. $L(\bar{\mathcal{A}}) = (2^P)^\omega \setminus L(\mathcal{A})$ and $|\bar{\mathcal{A}}| = |\mathcal{A}|$;
2. \mathcal{A}_1 s.t. $L(\mathcal{A}_1) = L(\mathcal{A}) \cup L(\mathcal{B})$ and $|\mathcal{A}_1| = \mathcal{O}(|\mathcal{A}| + |\mathcal{B}|)$;
3. \mathcal{A}_2 s.t. $L(\mathcal{A}_2) = L(\mathcal{A}) \cap L(\mathcal{B})$ and $|\mathcal{A}_2| = \mathcal{O}(|\mathcal{A}| + |\mathcal{B}|)$;

Proof. (1) Let $\mathcal{A} = (Q, AP, q_0, \delta, \Omega)$. We define $\bar{\mathcal{A}} = (Q, AP, q_0, \bar{\delta}, \bar{\Omega})$ where $\bar{\Omega}(q) := \Omega(q) + 1$, and $\bar{\delta}(q) := \delta(q)$ where $\bar{p} = \neg p$, $\overline{\neg p} = p$, $\overline{\zeta \vee \eta} = \bar{\zeta} \wedge \bar{\eta}$ and $\overline{\zeta \wedge \eta} = \bar{\zeta} \vee \bar{\eta}$. Clearly, the dual APA is not any bigger than its counterpart. A careful inspection reveals that it recognises the complement language.

(2+3) The APA \mathcal{A}_1 and \mathcal{A}_2 are obtained by taking the union of \mathcal{A} and \mathcal{B} and adding a new state q_0 with arbitrary priority and transitions obtained as the disjunction, resp. conjunction of the transitions of the two original starting states. \square

This makes APA a rich model for various decision problems.

3.4 Decision Problems for Alternating Parity Automata

Important decision problems for automata are the following.

- Non-Emptiness: given an APA \mathcal{A} , is $L(\mathcal{A}) \neq \emptyset$?
- Universality: given an APA \mathcal{A} , is $L(\mathcal{A}) = (2^P)^\omega$?
- Subsumption: given APA \mathcal{A} and \mathcal{B} , is $L(\mathcal{A}) \subseteq L(\mathcal{B})$?

- Equivalence: given APA \mathcal{A} and \mathcal{B} , is $L(\mathcal{A}) = L(\mathcal{B})$?

Thanks to Prop. 5 these problems are all interreducible in linear time. For instance, equivalence reduces to non-emptiness because $L(\mathcal{A}) \neq L(\mathcal{B})$ iff $L(\mathcal{C}) \neq \emptyset$ where \mathcal{C} recognises $(L(\mathcal{A}) \cap L(\overline{\mathcal{B}})) \cup (L(\overline{\mathcal{A}}) \cap L(\mathcal{B}))$. Equally, non-emptiness and universality are interreducible because $L(\mathcal{A}) \neq \emptyset$ iff $L(\overline{\mathcal{A}}) \neq (2^P)^\omega$.

Note that these reductions need the full power of alternation as well as the full power of the parity acceptance condition unless they are weak. The dual of a Büchi automaton for instance is not a Büchi automaton but a co-Büchi automaton and vice-versa. Weakness is preserved by dualisation though. Also, the dual of a nondeterministic automaton is in general not a nondeterministic automaton anymore. However, regarding it as an alternating automaton enables easy dualisation.

3.5 Alternating Parity Automata and Temporal Logics

Many decision problems for linear-time temporal logics can be phrased as an emptiness problem for (a subclass) of alternating parity automata. The crucial ingredient for this is of course an equivalence-preserving translation from formulas to automata. With Prop. 4 above it suffices to check that LT_μ can be translated into alternating parity automata that way. However, translating logics like LTL and PSL separately can be beneficial because they may not need the full power of the parity acceptance condition.

Proposition 6 ([23, 14, 10, 4]). *For every ...*

- LTL formula φ there is a weak APA \mathcal{A}_φ s.t. $L(\mathcal{A}_\varphi) = L(\varphi)$ and $|\mathcal{A}_\varphi| = \mathcal{O}(|\varphi|)$;
- PSL formula φ there is a weak APA \mathcal{A}_φ s.t. $L(\mathcal{A}_\varphi) = L(\varphi)$ and $|\mathcal{A}_\varphi| = 2^{\mathcal{O}(|\varphi|)}$;
- LT_μ formula φ there is an APA \mathcal{A}_φ s.t. $L(\mathcal{A}_\varphi) = L(\varphi)$ and $|\mathcal{A}_\varphi| = \mathcal{O}(|\varphi|)$.


4 The Size-Change Termination Principle for Alternating Parity Automata

4.1 Boxes and Their Composition

We use $<$ to denote the usual total ordering on \mathbb{N} or \mathbb{Z} , and introduce a second (non-well-founded) total ordering called *reward ordering*: $i \prec j$ iff $(-1)^i \cdot i < (-1)^j \cdot j$. I.e. we have $\dots \prec 3 \prec 1 \prec 0 \prec 2 \prec \dots$

For the remainder of this section we fix an APA $\mathcal{A} = (Q, \text{AP}, q_0, \delta, \Omega)$. Let $P = \{\Omega(q) \mid q \in Q\}$ be the set of all priorities occurring in \mathcal{A} . We write P_\perp for $P \cup \{\perp\}$ which will be used to model partial functions into P .


A *box* is an element of type $Q \times Q \rightarrow P_\perp$. The name suggest a particular visual representation of such a function. We regard the elements of Q as ports of a circuit, and a box has in- and out-ports which can be connected by an edge labeled with a number. For instance, if $Q = \{0, 1, 2\}$ and $P = \{3, 4, 5\}$ then the following is a box over Q .



We define a *composition* operation \circ on two boxes f, g by

$$(f \circ g)(q, p) = \max_{\prec} \{ \max_{\prec} \{ f(q, q'), g(q', p) \} \mid q' \in Q \}$$

Here we use the important convention that \perp is the *maximal* element w.r.t. \prec , but the *minimal* element w.r.t. \prec . Then box composition works as one would expect it from the graphical representation.



Box composition is lifted to sets of boxes in the natural way: $F \circ G := \{ f \circ g \mid f \in F, g \in G \}$.

With every finite word $v \in (2^{\text{AP}})^\omega$ we associate a set of boxes $[v]$ as follows: $[\epsilon]$ contains only a single box f which is defined by $f(q, q') = \Omega(q)$ if $q' = q$ and $f(q, q') = \perp$ otherwise. Furthermore, for every one-letter word $a \subseteq \text{AP}$ we have that $f \in [a]$ if for all $q, q' \in Q$:

1. $f(q, q') \in \{ \Omega(q), \perp \}$, and
2. $\{ p \mid f(q, p) \neq \perp \} \cup a \models \delta(q)$

Intuitively, a box belongs to $[a]$ if it connects every in-port to all the out-ports in some model that agrees with a on the atomic propositions and their negations. The labels on the connections simply reflect the priority of the in-port.

Using composition it is easy to define $[v]$ for longer words v : $[av] := [a] \circ [v]$. We write $[\mathcal{A}^*]$ for $\{ [v] \mid v \in (2^{\text{AP}})^* \}$ and $[\mathcal{A}^+]$ for $\{ [v] \mid v \in (2^{\text{AP}})^+ \}$. Note that these are finite sets.

A box f is called *idempotent* if $f \circ f = f$. It is called *good* w.r.t. some $Q' \subseteq Q$ if for all $q \in Q'$ we have that $f(q, q)$ is even.

4.2 Characterising the Emptiness Problem for Alternating Parity Automata

Theorem 1. $L(\mathcal{A}) \neq \emptyset$ iff there are $f \in [\mathcal{A}^*]$ and $g \in [\mathcal{A}^+]$ s.t. g is idempotent and good w.r.t. $\{ q \mid f(q_0, q) \neq \perp \}$.

Proof. “ \Leftarrow ” Suppose such $f, g \in [\mathcal{A}]$ exist. Then there must be words $v_f = a_1 \dots a_n$ and $v_g = b_1 \dots b_m$ with $n \geq 0$, $m \geq 1$ s.t. $f \in [v_f]$ and $g \in [v_g]$. I.e. there must be $f_1 \in [a_1], \dots, f_n \in [a_n], g_1 \in [b_1], \dots, g_m \in [b_m]$ s.t. $f =$

$f_1 \circ \dots \circ f_n$ and $g = g_1 \circ \dots \circ g_m$. A run on $v_f(v_g)^\omega$ can easily be extracted from the sequence $f_1, \dots, f_n, g_1, \dots, g_m, g_1, \dots$ by following all connections starting from q_0 . Suppose this run was not accepting. Then it would contain a path on which the highest priority seen infinitely often is odd. Note that idempotency of g means that this path is compressed into a connection from some q' to itself in this box. Furthermore, q' must be reachable from q_0 in f . But then $g(q', q')$ must be odd which contradicts the assumption that g is good.

“ \Rightarrow ” Suppose that $L(\mathcal{A}) \neq \emptyset$, i.e. there is a $w \in (2^{\text{AP}})^\omega$ s.t. $w \in L(\mathcal{A})$. Then there is an accepting run t of \mathcal{A} on w . Let $w = a_0 a_1 \dots$. The run t can easily be transformed into a sequence of boxes f_0, f_1, \dots by possibly adding nodes to each level s.t. every state is present on every level, and adding corresponding connections. Now consider a colouring of all ordered pairs of levels i, j with $i \leq j$, assigning to this pair the box $f_{i,j} := f_i \circ \dots \circ f_j$. Since there are only finitely many boxes, Ramsey’s Theorem [18] gives us an infinite sequence $j_0 < j_1 < \dots$ of indices s.t. all pairs of indices from this sequence get assigned to the same box. In particular we have $f_{j_0, j_1} = f_{j_1, j_2} = f_{j_0, j_2} = f_{j_0, j_1} \circ f_{j_1, j_2}$ which shows that it is idempotent. Define the required boxes as $f := f_{0, j_0}$ and $g := f_{j_0, j_1}$. What remains to be seen is that g is good w.r.t. the set of all states that q_0 is connected to in f . As above, suppose it was not. Then the run would have contained an infinite path on which the highest priority occurring infinitely often was odd which would contradict the assumption that it was accepting. \square

References

1. P. A. Abdulla, Y.-F. Chen, L. Clemente, L. Holík, C.-D. Hong, R. Mayr, and T. Vojnar. Simulation subsumption in ramsey-based Büchi automata universality and inclusion testing. In *Proc. 22nd Int. Conf. on Computer Aided Verification, CAV’10*, volume 6174 of *LNCS*, pages 132–147. Springer, 2010.
2. Inc. Accellera Organization. Formal semantics of Accellera property specification language, 2004. In Appendix B of <http://www.eda.org/vfv/docs/PSL-v1.1.pdf>.
3. B. Banieqbal and H. Barringer. Temporal logic with fixed points. In *Proc. Coll. on Temporal Logic in Specification*, volume 398 of *LNCS*, pages 62–73. Springer, 1989.
4. D. Bustan, D. Fisman, and J. Havlicek. Automata constructions for PSL. Technical Report MCS05-04, The Weizmann Institute of Science, 2005.
5. C. Dax, M. Hofmann, and M. Lange. A proof system for the linear time μ -calculus. In *Proc. 26th Conf. on Foundations of Software Technology and Theoretical Computer Science, FSTTCS’06*, volume 4337 of *LNCS*, pages 274–285. Springer, 2006.
6. C. Dax and F. Klaedtke. Alternation elimination by complementation (extended abstract). In *Proc. 15th Int. Conf. on Logic for Programming, Artificial Intelligence, and Reasoning, LPAR’08*, volume 5330 of *LNCS*, pages 214–229. Springer, 2008.
7. E. A. Emerson. Temporal and modal logic. In J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, volume B: Formal Models and Semantics, chapter 16, pages 996–1072. Elsevier and MIT Press, New York, USA, 1990.
8. S. Fogarty and M. Y. Vardi. Büchi complementation and size-change termination. In *Proc. 15th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems, TACAS’09*, volume 5505 of *LNCS*, pages 16–30. Springer, 2009.

9. S. Fogarty and M. Y. Vardi. Efficient Büchi universality checking. In *Proc. 16th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems, TACAS'10*, volume 6015 of *LNCS*, pages 205–220. Springer, 2010.
10. R. Kaivola. *Using Automata to Characterise Fixed Point Temporal Logics*. PhD thesis, LFCS, Division of Informatics, The University of Edinburgh, 1997. Tech. Rep. ECS-LFCS-97-356.
11. O. Kupferman, N. Piterman, and M. Y. Vardi. Extended temporal logic revisited. In *Proc. 12th Int. Conf. on Concurrency Theory, CONCUR'01*, volume 2154 of *LNCS*, pages 519–535. Springer, 2001.
12. M. Lange. Linear time logics around PSL: Complexity, expressiveness, and a little bit of succinctness. In *Proc. 18th Int. Conf. on Concurrency Theory, CONCUR'07*, volume 4703 of *LNCS*, pages 90–104. Springer, 2007.
13. C. S. Lee, N. D. Jones, and A. M. Ben-Amram. The size-change principle for program termination. *ACM SIGPLAN Notices*, 36(3):81–92, 2001.
14. C. Löding and W. Thomas. Alternating automata and logics over infinite words. In *Proc. IFIP Int. Conf. on Theoretical Computer Science, IFIP TCS2000*, volume 1872 of *LNCS*, pages 521–535. Springer, 2000.
15. S. Miyano and T. Hayashi. Alternating finite automata on omega-words. *TCS*, 32(3):321–330, 1984.
16. A. Pnueli. The temporal logic of programs. In *Proc. 18th Symp. on Foundations of Computer Science, FOCS'77*, pages 46–57, Providence, RI, USA, 1977. IEEE.
17. F. P. Ramsey. On a problem of formal logic. *Proc. London Mathematical Society, Series 2*, 30(4):338–384, 1928.
18. F. P. Ramsey. On a problem in formal logic. *Proc. London Math. Soc. (3)*, 30:264–286, 1930.
19. A. P. Sistla and E. M. Clarke. The complexity of propositional linear temporal logics. *Journal of the Association for Computing Machinery*, 32(3):733–749, 1985.
20. C. Stirling. Comparing linear and branching time temporal logics. In B. Ban-
ieqbal, H. Barringer, and A. Pnueli, editors, *Proc. Conf. on Temporal Logic in
Specification*, volume 398 of *LNCS*, pages 1–20, Berlin, 1989. Springer.
21. Vardi. Linear vs. branching time: A complexity-theoretic perspective. In *LICS:
IEEE Symposium on Logic in Computer Science*, 1998.
22. M. Y. Vardi. A temporal fixpoint calculus. In ACM, editor, *Proc. Conf. on Princi-
ples of Programming Languages, POPL'88*, pages 250–259, NY, USA, 1988. ACM
Press.
23. M. Y. Vardi. Alternating automata and program verification. *LNCS*, 1000:471–485,
1995.
24. M. Y. Vardi. *An Automata-Theoretic Approach to Linear Temporal Logic*, volume
1043 of *LNCS*, pages 238–266. Springer, New York, NY, USA, 1996.
25. M. Y. Vardi. Branching vs. linear time: Final showdown. In T. Margaria and
W. Yi, editors, *Proc. 7th Int. Conf. on Tools and Algorithms for the Construction
and Analysis of Systems, TACAS'01*, volume 2031 of *LNCS*, pages 1–22. Springer,
2001.
26. M. Y. Vardi and P. Wolper. Reasoning about infinite computations. *Information
and Computation*, 115(1):1–37, 1994.
27. P. Wolper. Temporal logic can be more expressive. *Information and Control*,
56:72–99, 1983.