An Efficient Normalisation Procedure for Linear Temporal Logic and Very Weak Alternating Automata

Salomon Sickert
Technische Universität München
Germany
s.sickert@tum.de

Javier Esparza
Technische Universität München
Germany
esparza@in.tum.de

Abstract

In the mid 80s, Lichtenstein, Pnueli, and Zuck proved a classical theorem stating that every formula of Past LTL (the extension of LTL with past operators) is equivalent to a formula of the form $\bigwedge_{i=1}^n \text{GF}\varphi_i \vee \text{FG}\psi_i$, where φ_i and ψ_i contain only past operators. Some years later, Chang, Manna, and Pnueli built on this result to derive a similar normal form for LTL. Both normalisation procedures have a non-elementary worst-case blow-up, and follow an involved path from formulas to counter-free automata to star-free regular expressions and back to formulas. We improve on both points. We present a direct and purely syntactic normalisation procedure for LTL yielding a normal form, comparable to the one by Chang, Manna, and Pnueli, that has only a single exponential blow-up. As an application, we derive a simple algorithm to translate LTL into deterministic Rabin automata. The algorithm normalises the formula, translates it into a special very weak alternating automaton, and applies a simple determinisation procedure, valid only for these special automata.

CCS Concepts: • Theory of computation $\rightarrow$ Modal and temporal logics; Automata over infinite objects.

Keywords: Linear Temporal Logic, Normal Form, Weak Alternating Automata, Deterministic Automata

ACM Reference Format:

Salomon Sickert and Javier Esparza. 2020. An Efficient Normalisation Procedure for Linear Temporal Logic and Very Weak Alternating Automata. In *Proceedings of the 35th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '20)*, July 8–11, 2020, Saarbrücken, Germany. ACM, New York, NY, USA, 14 pages. <https://doi.org/10.1145/3373718.3394743>

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

LICS '20, July 8–11, 2020, Saarbrücken, Germany

© 2020 Copyright held by the owner/author(s).

ACM ISBN 978-1-4503-7104-9/20/07.

<https://doi.org/10.1145/3373718.3394743>

1 Introduction

In seminal work carried out in the middle 80s, Lichtenstein, Pnueli, and Zuck investigated Past Linear Temporal Logic (Past LTL), a temporal logic with future and past operators. They proved the classical result stating that every formula is equivalent to another one of the form

$$\bigwedge_{i=1}^n \text{GF}\varphi_i \vee \text{FG}\psi_i \quad (1)$$

where φ_i and ψ_i only contain past operators [8, 25]. Shortly after, Manna and Pnueli introduced the *safety-progress* hierarchy, containing six classes of properties (Figure 1a), and presented a logical characterisation of each class in terms of syntactic fragments of Past LTL [11, 12]. The class of *reactivity properties*, placed at the top of the hierarchy, contains all Past LTL properties, and its syntactic characterisation, given by (1), is the class of *reactivity formulas*.

In the early 90s, LTL (which only has future operators, but is known to be as expressive as Past LTL), became the logic of choice for most model-checking applications. At that time Chang, Manna, and Pnueli showed that the classes of the safety-progress hierarchy also admit syntactic characterisations in terms of LTL fragments [3]. In particular, they proved that every LTL formula is equivalent to another one in which every path of the syntax tree alternates at most once between the “least-fixed-point” operators **U** and **M** and the “greatest-fixed-point” operators **W** and **R**. In the notation introduced in [2], which mimics the definition of the Σ_i , Π_i , and Δ_i classes of the arithmetical and polynomial hierarchies, they proved that every LTL formula is equivalent to a Δ_2 -formula.

While these normal forms have had large conceptual impact in model checking, automatic synthesis, and deductive verification (see e.g. [17] for a recent survey), the normalisation *procedures* proving that they are indeed normal forms have had none. In particular, contrary to the case of propositional or first-order logic, they have not been implemented in tools. The reason is that they are not direct, have high complexity, and their correctness proofs are involved. Let us elaborate on this. In [25], Zuck gives a detailed description of the normalisation procedure of [8]. First, Zuck translates the initial Past LTL formula into a counter-free semi-automaton, then applies the Krohn-Rhodes decomposition and other results to translate the automaton into a

star-free regular expression, and finally translates this expression into a reactivity formula with a non-elementary blow-up. In [11, 12] the procedure is not even presented, the reader is referred to [25] and/or to previous results¹. The normalisation procedure of [3] for LTL calls the translation procedure of [8, 25] for Past LTL as a subroutine, and so it is not any simpler². Finally, while Maler and Pnueli present in [10] an improved translation of star-free regular languages to Past LTL, their work still leads to a triple exponential normalisation procedure for Past LTL. Further, it is not clear to us if this translation can also be used to obtain Δ_2 -formulas.

In this paper we present a novel normalisation procedure that translates any LTL formula into an equivalent Δ_2 -formula. Our procedure is:

- *Direct*. It does not require any detour through automata or regular expressions.
- *Syntax-guided*. It consists of a few syntactic rewrite rules—not unlike the rules for putting a boolean formula in conjunctive or disjunctive normal form—that can be described in less than a page.
- *Single exponential*. The length of the Δ_2 -formula is at most exponential in the length of the original formula, a dramatic improvement on the previous non-elementary and triple exponential bounds.

The correctness proof of the procedure consists of a few lemmas, all of them with routine proofs by structural induction. It is presented in Sections 4 to 6, modulo the omission of some straightforward induction cases. To make this paper self-contained, the proofs of three lemmas taken from [5, 21] are reproduced in the appendix of the extended version of this paper [23]. We have mechanised the complete correctness proof in Isabelle/HOL [15], building upon previous work [1, 19, 20]. The formalised proof consists of roughly 1000 lines, from which one can extract a formally verified normalisation procedure consisting of ca. 200 lines of Standard ML code, excluding standard definitions added by the code-generation. Both the formalisation and instructions for extracting code are located in [22].

In the second part of the paper (Sections 7 and 8) we use the new normalisation procedure to derive a simple translation of LTL into deterministic Rabin automata (DRW). First, we show that every formula of Δ_2 can be translated into a very weak alternating Büchi automaton (A1W) in which every path has at most one alternation between accepting and non-accepting states. Further, we provide a simple determinisation procedure for these automata, based on a breakpoint construction. The LTL-to-DRW translation normalises the formula, transforms it into an A1W with at most one alternation, and determinises this intermediate automaton.

¹Including papers by Burgess, McNaughton and Pappet, Choueka, Thomas, and Gabby, Pnueli, Shela, and Stavi.

²Further, [3] only contains a short sketch of the translation of reactivity formulas into Δ_2 -formulas.

Due to space constraints we do not provide an overview of LTL-to-DRW translations and refer the reader to [21, Ch. 1]. Furthermore, we only provide a preliminary experimental evaluation of the proposed translations and leave a detailed analysis as future work.

2 Preliminaries

Let Σ be a finite alphabet. A *word* w over Σ is an infinite sequence of letters $a_0a_1a_2 \dots$ with $a_i \in \Sigma$ for all $i \geq 0$, and a *language* is a set of words. A *finite word* is a finite sequence of letters. As usual, the set of all words (finite words) is denoted Σ^ω (Σ^*). We let $w[i]$ (starting at $i = 0$) denote the i -th letter of a word w . The finite infix $w[i]w[i+1] \dots w[j-1]$ is abbreviated with w_{ij} and the infinite suffix $w[i]w[i+1] \dots$ with w_i . We denote the infinite repetition of a finite word $\sigma_1 \dots \sigma_n$ by $(\sigma_1 \dots \sigma_n)^\omega = \sigma_1 \dots \sigma_n \sigma_1 \dots \sigma_n \sigma_1 \dots$.

Definition 1. *LTL formulas over a set Ap of atomic propositions are constructed by the following syntax:*

$$\begin{aligned} \varphi ::= & \mathbf{tt} \mid \mathbf{ff} \mid a \mid \neg a \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \\ & \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi \mid \varphi \mathbf{W}\varphi \mid \varphi \mathbf{R}\varphi \mid \varphi \mathbf{M}\varphi \end{aligned}$$

where $a \in Ap$ is an atomic proposition and $\mathbf{X}$, $\mathbf{U}$, $\mathbf{W}$, $\mathbf{R}$, and $\mathbf{M}$ are the next, (strong) until, weak until, (weak) release, and strong release operators, respectively.

The inclusion of both the strong and weak until operators as well as the negation normal form are essential to our approach. The operators $\mathbf{R}$ and $\mathbf{M}$, however, are added to ensure that every formula of length n in the standard syntax, with negation but only the until operator, is equivalent to a formula of length $O(n)$ in our syntax. They could be removed, if we accept an exponential blow-up incurred by expressing $\mathbf{R}$ with $\mathbf{W}$. The semantics is defined as usual:

Definition 2. *Let w be a word over the alphabet 2^{Ap} and let φ be a formula. The satisfaction relation $w \models \varphi$ is inductively defined as the smallest relation satisfying:*

$$\begin{aligned} w &\models \mathbf{tt} \\ w &\models a \quad \text{iff } a \in w[0] \\ w &\models \neg a \quad \text{iff } a \notin w[0] \\ w &\models \varphi \wedge \psi \quad \text{iff } w \models \varphi \text{ and } w \models \psi \\ w &\models \varphi \vee \psi \quad \text{iff } w \models \varphi \text{ or } w \models \psi \\ w &\models \mathbf{X}\varphi \quad \text{iff } w_1 \models \varphi \\ w &\models \varphi \mathbf{U}\psi \quad \text{iff } \exists k. w_k \models \psi \text{ and } \forall j < k. w_j \models \varphi \\ w &\models \varphi \mathbf{M}\psi \quad \text{iff } \exists k. w_k \models \varphi \text{ and } \forall j \leq k. w_j \models \psi \\ w &\models \varphi \mathbf{R}\psi \quad \text{iff } \forall k. w_k \models \psi \text{ or } w \models \varphi \mathbf{M}\psi \\ w &\models \varphi \mathbf{W}\psi \quad \text{iff } \forall k. w_k \models \varphi \text{ or } w \models \varphi \mathbf{U}\psi \end{aligned}$$

We let $\mathcal{L}(\varphi) := \{w \in (2^{Ap})^\omega : w \models \varphi\}$ denote the language of φ . We overload the definition of $\models$ and write $\varphi \models \psi$ as a shorthand for $\mathcal{L}(\varphi) \subseteq \mathcal{L}(\psi)$.

We use the standard abbreviations $\mathbf{F}\varphi := \mathbf{tt} \mathbf{U} \varphi$ (eventually) and $\mathbf{G}\varphi := \mathbf{ff} \mathbf{R} \varphi$ (always). Finally, we introduce the

notion of equivalence of formulas, and equivalence within a language.

Definition 3. Two formulas φ and ψ are equivalent, denoted $\varphi \equiv \psi$, if $\mathcal{L}(\varphi) = \mathcal{L}(\psi)$. Given a language $L \subseteq (2^{Ap})^\omega$, two formulas φ and ψ are equivalent within L , denoted $\varphi \equiv^L \psi$, if $\mathcal{L}(\varphi) \cap L = \mathcal{L}(\psi) \cap L$.

3 The Safety-Progress Hierarchy

We recall the hierarchy of temporal properties studied by Manna and Pnueli [11] following the formulation of Černá and Pelánek [2]. In the ensuing sections we describe structures that have a direct correspondence to this hierarchy and in this sense the hierarchy provides a map to navigate the results of this paper.

Definition 4 ([2, 11]). Let $P \subseteq \Sigma^\omega$ be a property over Σ .

- P is a *safety property* if there exists a language of finite words $L \subseteq \Sigma^*$ such that for every $w \in P$ all finite prefixes of w belong to L .
- P is a *guarantee property* if there exists a language of finite words $L \subseteq \Sigma^*$ such that for every $w \in P$ there exists a finite prefix of w which belongs to L .
- P is an *obligation property* if it can be expressed as a positive boolean combination of safety and guarantee properties.
- P is a *recurrence property* if there exists a language of finite words $L \subseteq \Sigma^*$ such that for every $w \in P$ infinitely many prefixes of w belong to L .
- P is a *persistence property* if there exists a language of finite words $L \subseteq \Sigma^*$ such that for every $w \in P$ all but finitely many prefixes of w belong to L .
- P is a *reactivity property* if P can be expressed as a positive boolean combination of recurrence and persistence properties.

The inclusions between these classes are shown in Figure 1a. Chang, Manna, and Pnueli give in [3] a syntactic characterisation of the classes of the safety-progress hierarchy in terms of fragments of LTL. The following is a corollary of the proof of [3, Thm. 8]:

Definition 5 (Adapted from [2]). We define the following classes of LTL formulas:

- The class $\Sigma_0 = \Pi_0 = \Delta_0$ is the least set containing all atomic propositions and their negations, and is closed under the application of conjunction and disjunction.
- The class Σ_{i+1} is the least set containing Π_i and is closed under the application of conjunction, disjunction, and the X , U , and M operators.
- The class Π_{i+1} is the least set containing Σ_i and is closed under the application of conjunction, disjunction, and the X , R , and W operators.
- The class Δ_{i+1} is the least set containing Σ_{i+1} and Π_{i+1} and is closed under the application of conjunction and disjunction.

(a) Safety-progress hierarchy [11] (b) Syntactic-future hierarchy

Figure 1. Both hierarchies, side-by-side, indicating the correspondence of Theorem 6

Theorem 6 (Adapted from [2]). A property that is specifiable in LTL is a guarantee (safety, obligation, persistence, recurrence, reactivity, respectively) property if and only if it is specifiable by a formula from the class $\Sigma_1, (\Pi_1, \Delta_1, \Sigma_2, \Pi_2, \Delta_2, \text{respectively})$.

4 Overview of the Normalisation Result

Fix an LTL formula φ over a set of atomic propositions Ap . Our new normal form is based on two notions:

- A partition of the universe $\mathcal{U} := (2^{Ap})^\omega$ of all words into equivalence classes of words that, loosely speaking, exhibit the same “limit-behaviour” with respect to φ .
- The notion of stable word with respect to φ .

A partition of $\mathcal{U}$. Let $\mu(\varphi)$ and $\nu(\varphi)$ be the sets containing the subformulas of φ of the form $\psi_1 op \psi_2$ for $op \in \{U, M\}$ and $op \in \{W, R\}$, respectively. Given a word w , define:

$$\begin{aligned} \mathcal{GF}_w^\varphi &:= \{\psi : \psi \in \mu(\varphi) \wedge w \models \mathbf{GF}\psi\} \\ \mathcal{FG}_w^\varphi &:= \{\psi : \psi \in \nu(\varphi) \wedge w \models \mathbf{FG}\psi\} \end{aligned}$$

(To simplify the notation, when φ is clear from the context we simply write $\mathcal{GF}_w$ and $\mathcal{FG}_w$.) Two words w, v have the same *limit-behaviour* w.r.t. φ if $\mathcal{GF}_w = \mathcal{GF}_v$ and $\mathcal{FG}_w = \mathcal{FG}_v$. Having the same limit-behaviour is an equivalence relation, which induces the partition $\mathcal{P} = \{\mathcal{P}_{M,N} \subseteq \mathcal{U} : M \subseteq \mu(\varphi), N \subseteq \nu(\varphi)\}$ given by:

$$\mathcal{P}_{M,N} := \{w \in \mathcal{U} : M = \mathcal{GF}_w \wedge N = \mathcal{FG}_w\} \quad (2)$$

Example 7. Let $\varphi = Ga \vee bUc$. We have $\mu(\varphi) = \{bUc\}$ and $\nu(\varphi) = \{Ga\}$. The partition $\mathcal{P}$ has four equivalence classes:

- $\mathcal{P}_{\emptyset, \emptyset}$ contains all words such that bUc holds only finitely often and Ga fails infinitely often (which in this case implies that Ga never holds), e.g. $\{b\}^\omega$ or $\{c\}\{b\}^\omega$.
- $\mathcal{P}_{\emptyset, \{Ga\}}$ contains all words such that bUc holds finitely often and Ga fails finitely often, e.g. $\{a\}^\omega$ or $\{c\}\{a\}^\omega$.
- $\mathcal{P}_{\{bUc\}, \emptyset}$ contains all words such that bUc holds infinitely often and Ga fails infinitely often, e.g. $(\{a\}\{c\})^\omega$ or $\{a\}\{c\}^\omega$.

$\mathcal{P}_{\emptyset, \emptyset}$		$\mathcal{P}_{\{bUc\}, \emptyset}$	
$\{c\}\{b\}^\omega$	$\{b\}^\omega$	$(\{a\}\{c\})^\omega$	$\{a\}\{c\}^\omega$
$\{c\}\{a\}^\omega$	$\{a\}^\omega$	$(\{a, c\}\{a\})^\omega$	$\{b\}\{a, c\}^\omega$
$\mathcal{P}_{\emptyset, \{Ga\}}$		$\mathcal{P}_{\{bUc\}, \{Ga\}}$	

Figure 2. Partition of $(2^{\{a,b,c\}})^\omega$ according to $\varphi = Ga \vee bUc$.

- $\mathcal{P}_{\{bUc\}, \{Ga\}}$ contains all words such that bUc holds infinitely often and Ga fails finitely often, e.g. $\{b\}\{a, c\}^\omega$ or $(\{a, c\}\{a\})^\omega$.

The partition is graphically shown in Figure 2. The equivalence classes are shown in blue, red, yellow, and green (ignore the inner part in darker colour for the moment).

Stable words. A word w is *stable* with respect to φ if every formula of $\mu(\varphi)$ holds either never or infinitely often along w (i.e., either none or infinitely many of its suffixes satisfy the formula), and every formula of $\nu(\varphi)$ fails never or infinitely often along w . In particular, for a stable word no formula of $\mu(\varphi)$ can hold a finite, nonzero number of times before it fails forever, and no formula of $\nu(\varphi)$ can fail a finite, nonzero number of times before it holds forever. It follows immediately from this definition that not every word is stable, but every word eventually stabilises, meaning that all but finitely many of its suffixes are stable. Let S_φ denote the set of stable words with respect to φ . Defining

$$\begin{aligned}\mathcal{F}_w^\varphi &:= \{\psi : \psi \in \mu(\varphi) \wedge w \models F\psi\} \\ \mathcal{G}_w^\varphi &:= \{\psi : \psi \in \nu(\varphi) \wedge w \models G\psi\}\end{aligned}$$

we easily obtain:

$$S_\varphi := \{w \in \mathcal{U} : \mathcal{F}_w^\varphi = \mathcal{GF}_w^\varphi \wedge \mathcal{G}_w^\varphi = \mathcal{FG}_w^\varphi\} \quad (3)$$

Example 8. Let $\varphi = Ga \vee bUc$. The words $\{c\}^n \{a\}^\omega$ for $n \geq 1$ are not stable w.r.t. φ , because bUc holds exactly n times along the word. However, the suffix $\{a\}^\omega$ is stable. Figure 2 represents the stable words of each element $\mathcal{P}_{M,N}$ of the partition in darker colour, and gives examples of stable words for each class.

The starting point of this paper is the observation that some results of [5, 21] allow us to easily derive a normal form for LTL, albeit only when LTL is interpreted on stable words. More precisely, in Section 5 we show that for every $M \subseteq \mu(\varphi)$ and $N \subseteq \nu(\varphi)$ there exist formulas $\varphi[M]_1^\Pi \in \Pi_1$ and $\varphi[N]_1^\Sigma \in \Sigma_1$ such that:

$$\varphi \equiv^{S_\varphi} \bigvee_{\substack{M \subseteq \mu(\varphi) \\ N \subseteq \nu(\varphi)}} \left(\varphi[M]_1^\Pi \wedge \bigwedge_{\psi \in M} \text{GF}(\psi[N]_1^\Sigma) \wedge \bigwedge_{\psi \in N} \text{FG}(\psi[M]_1^\Pi) \right) \quad (4)$$

Further, $\varphi[M]_1^\Pi$ and $\varphi[N]_1^\Sigma$ are obtained from φ , M , and N by means of a simple, linear-time syntactic substitution procedure. Observe that the right-hand side is a formula of Δ_2 , and that we write $\equiv^{S_\varphi}$, i.e., the equivalence is only valid within the universe of stable words. In this paper we lift this restriction. In Section 6 we define a formula $\varphi[M]_2^\Sigma \in \Sigma_2$ by means of another linear-time, syntactic substitution procedure, such that:

$$\varphi \equiv \bigvee_{\substack{M \subseteq \mu(\varphi) \\ N \subseteq \nu(\varphi)}} \left(\varphi[M]_2^\Sigma \wedge \bigwedge_{\psi \in M} \text{GF}(\psi[N]_1^\Sigma) \wedge \bigwedge_{\psi \in N} \text{FG}(\psi[M]_1^\Pi) \right) \quad (5)$$

Example 9. For $\varphi = F(a \wedge G(b \vee Fc)) \in \Sigma_3$, the still-to-be-defined normal form (4) will yield:

$$\varphi \equiv^{S_\varphi} (\text{GF}a \wedge \text{FG}b) \vee (\text{GF}a \wedge \text{GF}c)$$

Indeed, since $\varphi \in \mu(\varphi)$, every stable word satisfying φ must satisfy it infinitely often, and so equivalence for stable words holds, although the formulas are not equivalent. For Equation (5) we will obtain:

$$\varphi \equiv F(a \wedge ((b \vee Fc) \text{ U } Gb)) \vee (Fa \wedge \text{GF}c)$$

Observe that the right-hand-side belongs to Δ_2 .

5 The Formulas $\varphi[M]_1^\Pi$ and $\varphi[N]_1^\Sigma$

We recall the definitions of the formulas $\varphi[M]_1^\Pi$ and $\varphi[N]_1^\Sigma$, introduced in [5, 21] with a slightly different notation.

The formula $\varphi[M]_1^\Pi$. Define $\mathcal{P}_M := \bigcup_{N \subseteq \nu(\varphi)} \mathcal{P}_{M,N}$. Observe that $\mathcal{P}_M$ is the language of the words w such that $M = \mathcal{GF}_w$. The formula $\varphi[M]_1^\Pi$ is defined with the goal of satisfying the following identity:

$$\varphi \equiv^{S_\varphi \cap \mathcal{P}_M} \varphi[M]_1^\Pi \quad (6)$$

Intuitively, the identity states that within the universe of the stable words of $\mathcal{P}_M$, the formula φ can be replaced by the simpler formula $\varphi[M]_1^\Pi$.

All insights required to define $\varphi[M]_1^\Pi$ are illustrated by the following examples, where we assume that $w \in S_\varphi \cap \mathcal{P}_M$:

- $\varphi = Fa \wedge Gb$ and $M = \{Fa\}$. Since $M = \mathcal{GF}_w$, we have $Fa \in \mathcal{GF}_w$, which implies $w \models \text{GF}a$. So $w \models Fa \wedge Gb$ iff $w \models Gb$, and so we can set $\varphi[M]_1^\Pi := \text{tt} \wedge Gb$, i.e., we can define $\varphi[M]_1^\Pi$ as the result of substituting tt for Fa in φ . The yet-to-be-defined substitution in-fact replaces the abbreviation $Fa = \text{tt}Ua$ by $\text{tt}Wa \equiv \text{tt}$.
- $\varphi = Fa \wedge Gb$ and $M = \emptyset$. Since $M = \mathcal{F}_w$, we have $Fa \notin \mathcal{F}_w$, and so $w \not\models Fa$. In other words, $w \models Fa \wedge Gb$ iff $w \models \text{ff}$, and so we can set $\varphi[M]_1^\Pi := \text{ff} \wedge Gb$.
- $\varphi = G(bUc)$ and $M = \{bUc\}$. Since $M = \mathcal{GF}_w$, we have $bUc \in \mathcal{GF}_w$, and so $w \models \text{GF}(bUc)$. This does not imply $w_i \models bUc$ for all suffixes of w , but it implies that c will hold infinitely often in the future. So

$w \models G(bUc)$ iff $w \models G(bWc)$, and so we can define $\varphi[M]_1^\Pi := G(bWc)$.

Definition 10 ([5, 21]). Let $M \subseteq \mu(\varphi)$ be a set of formulas. The formula $\varphi[M]_1^\Pi$ is inductively defined as follows:

$$\begin{aligned} (\varphi U \psi)[M]_1^\Pi &:= \begin{cases} \varphi[M]_1^\Pi \mathbf{W} \psi[M]_1^\Pi & \text{if } \varphi U \psi \in M \\ \mathbf{ff} & \text{otherwise.} \end{cases} \\ (\varphi \mathbf{M} \psi)[M]_1^\Pi &:= \begin{cases} \varphi[M]_1^\Pi \mathbf{R} \psi[M]_1^\Pi & \text{if } \varphi \mathbf{M} \psi \in M \\ \mathbf{ff} & \text{otherwise.} \end{cases} \end{aligned}$$

All other cases are defined homomorphically, e.g., $a[M]_1^\Pi := a$ for every $a \in \mathcal{A}$, $(X\varphi)[M]_1^\Pi := X(\varphi[M]_1^\Pi)$, and $(\varphi \mathbf{W} \psi)[M]_1^\Pi := (\varphi[M]_1^\Pi) \mathbf{W} (\psi[M]_1^\Pi)$.

The following lemma, proved in [5, 21], shows that $\varphi[M]_1^\Pi$ indeed satisfies Equation (6). Since the notation of [5, 21] is slightly different, we include proofs with the new notation for the cited results in the appendix of the extended version of this paper [23] for convenience.

Lemma 11 ([5, 21]). Let w be a word, and let $M \subseteq \mu(\varphi)$ be a set of formulas.

1. If $\mathcal{F}_w^\varphi \subseteq M$ and $w \models \varphi$, then $w \models \varphi[M]_1^\Pi$.
2. If $M \subseteq \mathcal{G}\mathcal{F}_w^\varphi$ and $w \models \varphi[M]_1^\Pi$, then $w \models \varphi$.
3. $\varphi \equiv_{S_\varphi \cap \mathcal{P}_M} \varphi[M]_1^\Pi$

Observe that the first two statements do not assume that w is stable. This is an aspect we will later make use of for the definition of the normalisation procedure.

The formula $\varphi[N]_1^\Sigma$. Let $\mathcal{P}_N := \bigcup_{M \subseteq \mu(\varphi)} \mathcal{P}_{M,N}$. The formula $\varphi[N]_1^\Sigma$ is designed to satisfy

$$\varphi \equiv_{S_\varphi \cap \mathcal{P}_N} \varphi[N]_1^\Sigma \quad (7)$$

and its definition is completely dual to that of $\varphi[M]_1^\Pi$.

Definition 12 ([5, 21]). Let $N \subseteq \nu(\varphi)$ be a set of formulas. The formula $\varphi[N]_1^\Sigma$ is inductively defined as follows:

$$\begin{aligned} (\varphi \mathbf{R} \psi)[N]_1^\Sigma &:= \begin{cases} \mathbf{tt} & \text{if } \varphi \mathbf{R} \psi \in N \\ \varphi[N]_1^\Sigma \mathbf{M} \psi[N]_1^\Sigma & \text{otherwise.} \end{cases} \\ (\varphi \mathbf{W} \psi)[N]_1^\Sigma &:= \begin{cases} \mathbf{tt} & \text{if } \varphi \mathbf{W} \psi \in N \\ \varphi[N]_1^\Sigma \mathbf{U} \psi[N]_1^\Sigma & \text{otherwise.} \end{cases} \end{aligned}$$

All other cases are defined homomorphically.

The dual of Lemma 11 also holds:

Lemma 13 ([5, 21]). Let w be a word, and let $N \subseteq \nu(\varphi)$ be a set of formulas.

1. If $\mathcal{F}\mathcal{G}_w^\varphi \subseteq N$ and $w \models \varphi$, then $w \models \varphi[N]_1^\Sigma$.
2. If $N \subseteq \mathcal{G}\mathcal{F}_w^\varphi$ and $w \models \varphi[N]_1^\Sigma$, then $w \models \varphi$.
3. $\varphi \equiv_{S_\varphi \cap \mathcal{P}_N} \varphi[N]_1^\Sigma$

A normal form for stable words. We use the following result from [5, 21] to characterise the stable words of a partition $\mathcal{P}_{M,N}$ that satisfy φ :

Lemma 14 ([5, 21]). Let w be a word, and let $M \subseteq \mu(\varphi)$ and $N \subseteq \nu(\varphi)$. Then define:

$$\Phi(M, N) := \bigwedge_{\psi \in M} \mathbf{GF}(\psi[N]_1^\Sigma) \wedge \bigwedge_{\psi \in N} \mathbf{FG}(\psi[M]_1^\Pi)$$

We have:

1. If $M = \mathcal{G}\mathcal{F}_w$ and $N = \mathcal{F}\mathcal{G}_w$, then $w \models \Phi(M, N)$.
2. If $w \models \Phi(M, N)$, then $M \subseteq \mathcal{G}\mathcal{F}_w$ and $N \subseteq \mathcal{F}\mathcal{G}_w$.

Equipped with this lemma, let us show that a stable word of $\mathcal{P}_{M,N}$ satisfies φ iff it satisfies $\varphi[M]_1^\Pi \wedge \Phi(M, N)$. Let w be a stable word of $\mathcal{P}_{M,N}$. If w satisfies φ , then it satisfies $\varphi[M]_1^\Pi$ by Lemma 11.3 and $\Phi(M, N)$ by Lemma 14.1 (recall that, since $w \in \mathcal{P}_{M,N}$, we have $M = \mathcal{G}\mathcal{F}_w$ and $N = \mathcal{F}\mathcal{G}_w$ by Equation (2)). For the other direction, assume that w satisfies $\varphi[M]_1^\Pi \wedge \Phi(M, N)$. Then we have $M \subseteq \mathcal{G}\mathcal{F}_w$ by Lemma 14.2 and so w satisfies φ by Lemma 11.2. (This direction does not even require stability.)

Since every word belongs to some element of the partition, we obtain a normal form for stable words:

Proposition 15.

$$\varphi \equiv_{S_\varphi} \bigvee_{\substack{M \subseteq \mu(\varphi) \\ N \subseteq \nu(\varphi)}} \left(\varphi[M]_1^\Pi \wedge \bigwedge_{\psi \in M} \mathbf{GF}(\psi[N]_1^\Sigma) \wedge \bigwedge_{\psi \in N} \mathbf{FG}(\psi[M]_1^\Pi) \right)$$

Proof. Define $\Phi(M, N)$ as in Lemma 14 and let $w \in S_\varphi$ be a stable word. We show that w satisfies φ iff it satisfies $\varphi[M]_1^\Pi$ and $\Phi(M, N)$ for some $M \subseteq \mu(\varphi)$ and $N \subseteq \nu(\varphi)$.

Assume $w \models \varphi$. Let $M := \mathcal{G}\mathcal{F}_w$ and $N := \mathcal{F}\mathcal{G}_w$. By Lemma 14.1 $w \models \Phi(M, N)$ holds. Since w is stable, we have $\mathcal{F}_w = \mathcal{G}\mathcal{F}_w = M$ (see Equation (3)). By Lemma 11.1 we have $w \models \varphi[M]_1^\Pi$, and we are done.

Assume $w \models (\varphi[M]_1^\Pi \wedge \Phi(M, N))$ for some $M \subseteq \mu(\varphi)$ and $N \subseteq \nu(\varphi)$. Using the second part of Lemma 14 we get $M \subseteq \mathcal{G}\mathcal{F}_w$. Applying Lemma 11.2 we get $w \models \varphi$. $\square$

Example 16. Let $\varphi = F(a \wedge G(b \vee Fc))$. We have $\mu(\varphi) = \{\varphi, Fc\}$ and $\nu(\varphi) = \{G(b \vee Fc)\}$. So there are four possible choices for M , and two for N . It follows that the right-hand-side of Proposition 15 has eight disjuncts. However, all disjuncts with $\varphi \notin M$ are equivalent to $\mathbf{ff}$ because then $\varphi[M]_1^\Pi = \mathbf{ff}$, and the same holds for all disjuncts with $\varphi \in M$ and $N = \emptyset$ because $\varphi[\emptyset]_1^\Sigma = \mathbf{ff}$.

The two remaining disjuncts are $M_1 = \{\varphi\}$, $N_1 = \{G(b \vee Fc)\}$, and $M_2 = \{\varphi, Fc\}$, $N_2 = \{G(b \vee Fc)\}$. For both we have $\varphi[M_1]_1^\Pi \equiv \varphi[M_2]_1^\Pi \equiv \mathbf{tt}$. Further, for the first disjunct we have

$$\mathbf{GF}(\varphi[N_1]_1^\Sigma) \wedge \mathbf{FG}((G(b \vee Fc))[M_1]_1^\Pi) \equiv \mathbf{GF}a \wedge \mathbf{FG}b$$

and for the second we get

$$\begin{aligned} & \mathbf{GF}(\varphi[N_2]_1^\Sigma) \wedge \mathbf{GF}(\mathbf{Fc})[N_2]_1^\Sigma \wedge \mathbf{FG}((\mathbf{G}(b \vee \mathbf{Fc})))[M_2]_1^\Pi) \\ & \equiv \mathbf{GF}a \wedge \mathbf{GF}c \wedge \mathbf{FG}(\mathbf{Gtt}) \equiv \mathbf{GF}a \wedge \mathbf{GF}c. \end{aligned}$$

Together we obtain $\mathbf{F}(a \wedge \mathbf{G}(b \vee \mathbf{Fc})) \equiv^{S_\varphi} \mathbf{GF}a \wedge (\mathbf{FG}b \vee \mathbf{GF}c)$.

6 A Normal Form for LTL

Proposition 15 has little interest in itself because of the restriction to stable words. However, it serves as the starting point for our search for an unrestricted normal form, valid for all words. Observe that Lemma 14 does not depend on w being stable. Contrary, Lemma 11.1 refers to $\mathcal{F}_w$ and we crucially depend on stability to replace it by $\mathcal{GF}_w$. Consequently, we only need to find a replacement for the first conjunct and can leave the rest of the structure, i.e. the enumeration of all possible combinations of $M \subseteq \mu(\varphi)$ and $\Phi(M, N)$, unchanged. More precisely, we search for a mapping $\varphi\langle\cdot\rangle$ that assigns to every $M \subseteq \mu(\varphi)$ a formula $\varphi\langle M \rangle \in \Sigma_2$ such that:

$$\varphi \equiv \bigvee_{\substack{M \subseteq \mu(\varphi) \\ N \subseteq \nu(\varphi)}} \left(\varphi\langle M \rangle \wedge \bigwedge_{\psi \in M} \mathbf{GF}(\psi[N]_1^\Sigma) \wedge \bigwedge_{\psi \in N} \mathbf{FG}(\psi[M]_1^\Pi) \right) \quad (8)$$

The following lemma gives sufficient conditions for $\varphi\langle M \rangle$.

Lemma 17. *For every $M \subseteq \mu(\varphi)$, let $\varphi\langle M \rangle$ be a formula satisfying:*

- (a) *For every $M' \subseteq \mu(\varphi)$: $M \subseteq M' \implies \varphi\langle M \rangle \models \varphi\langle M' \rangle$*
- (b) *For every word w : $w \models \varphi \iff w \models \varphi\langle \mathcal{GF}_w \rangle$*

Then Equation (8) holds.

Proof. Assume that (a, b) hold, and let w be a word. We show that w satisfies φ iff it satisfies the right-hand-side of (8).

($\implies$) Assume w satisfies φ . By (b) we have $w \models \varphi\langle \mathcal{GF}_w \rangle$. We claim that the disjunct of the right-hand-side of Equation (8) with $M := \mathcal{GF}_w$ and $N := \mathcal{F}\mathcal{G}_w$ holds. Indeed, $w \models \varphi\langle M \rangle$ trivially holds, and the rest follows from Lemma 14.1.

($\impliedby$) Assume w satisfies the right-hand side of Equation (8). Then there exist $M \subseteq \mu(\varphi)$ and $N \subseteq \nu(\varphi)$ such that $w \models \varphi\langle M \rangle$ holds, $w \models \mathbf{GF}(\psi[N]_1^\Sigma)$ holds for every $\psi \in M$, and $w \models \mathbf{FG}(\psi[M]_1^\Pi)$ holds for every $\psi \in N$. Lemma 14.2 yields $M \subseteq \mathcal{GF}_w$, and (a) yields $\varphi\langle \mathcal{GF}_w \rangle$. Applying (b) we get $w \models \varphi$. $\square$

Note that Lemma 17 can also be dualised and we could search for a mapping $\varphi\langle\cdot\rangle$ that assigns to every $N \subseteq \nu(\varphi)$ a formula $\varphi\langle N \rangle \in \Pi_2$ such that Equation (8) holds.

Unfortunately we cannot simply take $\varphi\langle M \rangle := \varphi[M]_1^\Pi$ or $\varphi\langle N \rangle := \varphi[N]_1^\Sigma$: Both choices satisfy condition (a) of Lemma 17, as proven by Lemma 18³, but fail to satisfy condition (b) as shown by Example 19.

Lemma 18. *$\varphi[\cdot]_1^\Pi$ and $\varphi[\cdot]_1^\Sigma$ have the following properties: For every $M, M' \subseteq \mu(\varphi)$ and $N, N' \subseteq \nu(\varphi)$:*

$$\begin{aligned} M \subseteq M' & \implies \varphi[M]_1^\Pi \models \varphi[M']_1^\Pi \\ N \subseteq N' & \implies \varphi[N]_1^\Sigma \models \varphi[N']_1^\Sigma \end{aligned}$$

Proof. (a) By induction on φ . We show only two cases, since all other cases are either trivial or analogous.

Case $\varphi = \psi_1 \mathbf{U} \psi_2$. Assume $w \models \varphi[M]_1^\Pi$ holds. Due to the definition of $\varphi[M]_1^\Pi$ we have $\varphi \in M$ and thus also $\varphi \in M'$. Thus we have $w \models (\psi_1[M]_1^\Pi) \mathbf{W}(\psi_2[M]_1^\Pi)$ and applying the induction hypothesis we get $w \models (\psi_1[M']_1^\Pi) \mathbf{W}(\psi_2[M']_1^\Pi)$. Hence $w \models \varphi[M']_1^\Pi$.

Case $\varphi = \psi_1 \mathbf{W} \psi_2$. Assume $w \models \varphi[N]_1^\Sigma$ holds. If $\varphi \in N'$ then $w \models \varphi[N']_1^\Sigma$ trivially holds. If $\varphi \notin N'$ then also $\varphi \notin N$, and we get $w \models (\psi_1[N]_1^\Sigma) \mathbf{U}(\psi_2[N]_1^\Sigma)$. Using the induction hypothesis we get $w \models (\psi_1[N']_1^\Sigma) \mathbf{U}(\psi_2[N']_1^\Sigma)$, and we are done. $\square$

Example 19. *Let us first exhibit a formula φ and a word w such that $w \models \varphi$, but $w \not\models \varphi[\mathcal{GF}_w]_1^\Pi$. For this take $\varphi = \mathbf{F}a$ and $w = \{a\}^\omega$. Thus $w \models \varphi$ and $\mathcal{GF}_w = \emptyset$. However, $(\mathbf{F}a)[\emptyset]_1^\Pi = \mathbf{ff}$ and hence $w \not\models (\mathbf{F}a)[\mathcal{GF}_w]_1^\Pi$.*

We now move to the second case. Let us exhibit φ and w such that $w \not\models \varphi$ and $w \models \varphi[\mathcal{F}\mathcal{G}_w]_1^\Sigma$. Dually, let $\varphi = \mathbf{G}a$ and $w = \{a\}^\omega$. Then $w \not\models \varphi$, but $\mathcal{F}\mathcal{G}_w = \{\mathbf{G}a\}$ and $(\mathbf{G}a)[\{\mathbf{G}a\}]_1^\Sigma = \mathbf{tt}$ and hence $w \models (\mathbf{G}a)[\mathcal{F}\mathcal{G}_w]_1^\Sigma$.

The key to finding a mapping $\varphi\langle\cdot\rangle$ satisfying both conditions of Lemma 17 is the technical result below, for which we offer the following intuition. The following equivalence is a valid law of LTL:

$$\mathbf{G}\varphi \equiv \varphi \mathbf{U} \mathbf{G}\varphi \quad (9)$$

In order to prove that a word w satisfies the right-hand-side we can take an *arbitrary* index $i \geq 0$, prove that $w_j \models \varphi$ holds for every $j < i$, and then prove that $w_i \models \mathbf{G}\varphi$. Since we are free to choose i , we can pick it such that w_i is a stable word, which allows us to apply the machinery of Section 5 and obtain:

Lemma 20. *For every word w :*

$$w \models \mathbf{G}\varphi \iff w \models \varphi \mathbf{U} \mathbf{G}(\varphi[\mathcal{GF}_w]_1^\Pi)$$

Proof. We prove both directions separately.

($\implies$) Assume $w \models \mathbf{G}\varphi$ holds. Let w_i be a stable suffix of w . By the definition of stability we have $\mathcal{F}_{w_i}^\varphi = \mathcal{F}_{w_j}^\varphi = \mathcal{GF}_w^\varphi$ for every $j \geq i$. By Lemma 11.1, we have

$$w_j \models \varphi \implies w_j \models \varphi[\mathcal{GF}_w^\varphi]_1^\Pi \text{ for every } j \geq i$$

and so in particular $w_i \models \mathbf{G}(\varphi[\mathcal{GF}_w^\varphi]_1^\Pi)$. We proceed as follows:

$$\begin{aligned} & w \models \mathbf{G}\varphi \\ \implies & w_i \models \mathbf{G}(\varphi[\mathcal{GF}_w^\varphi]_1^\Pi) \wedge \forall k < i. w_k \models \varphi \\ \implies & w \models \varphi \mathbf{U} \mathbf{G}(\varphi[\mathcal{GF}_w^\varphi]_1^\Pi) \end{aligned}$$

³This lemma is needed again for the proof of Theorem 23.

($\Leftarrow$) This is an immediate consequence of Lemma 11.2. $\square$

With the help of the standard LTL-equivalences

$$\varphi \mathbf{W} \psi \equiv \varphi \mathbf{U} (\psi \vee \mathbf{G} \varphi) \quad (10)$$

$$\varphi \mathbf{R} \psi \equiv (\varphi \vee \mathbf{G} \psi) \mathbf{M} \psi \quad (11)$$

Lemma 20 can be extended to a more powerful proposition.

Proposition 21. *For all formulas φ, ψ , and for every word w :*

$$\begin{aligned} w \models \varphi \mathbf{W} \psi &\iff w \models \varphi \mathbf{U} (\psi \vee \mathbf{G}(\varphi[\mathcal{GF}_w^\varphi]_1^\Pi)) \\ w \models \varphi \mathbf{R} \psi &\iff w \models (\varphi \vee \mathbf{G}(\psi[\mathcal{GF}_w^\psi]_1^\Pi)) \mathbf{M} \psi \end{aligned}$$

Proof. We only prove the first statement. The proof of the second is dual.

($\Rightarrow$) Assume $w \models \varphi \mathbf{W} \psi$. We split this branch of the proof further, by a case distinction on whether $w \models \mathbf{G} \varphi$ holds. If $w \models \mathbf{G} \varphi$ holds, then by Lemma 20 we have $w \models \varphi \mathbf{U} \mathbf{G}(\varphi[\mathcal{GF}_w^\varphi]_1^\Pi)$, and so $w \models \varphi \mathbf{U} (\psi \vee \mathbf{G}(\varphi[\mathcal{GF}_w^\varphi]_1^\Pi))$ holds. Assume now that $w \not\models \mathbf{G} \varphi$. Then we simply derive:

$$\begin{aligned} w &\models \varphi \mathbf{W} \psi \\ \iff w &\models \varphi \mathbf{U} \psi \quad (w \not\models \mathbf{G} \varphi) \\ \implies w &\models \varphi \mathbf{U} (\psi \vee \mathbf{G}(\varphi[\mathcal{GF}_w^\varphi]_1^\Pi)) \end{aligned}$$

($\Leftarrow$) By Lemma 11.2 we have $(w_j \models \varphi[\mathcal{GF}_w^\varphi]_1^\Pi \implies w_j \models \varphi)$ for all $j \geq 0$. Thus $w_j \models (\mathbf{G} \varphi)[\mathcal{GF}_w^\varphi]_1^\Pi \implies w_j \models \mathbf{G} \varphi$ for all $j \geq 0$ and we can simply derive:

$$\begin{aligned} w &\models \varphi \mathbf{U} (\psi \vee \mathbf{G}(\varphi[\mathcal{GF}_w^\varphi]_1^\Pi)) \\ \implies w &\models \varphi \mathbf{U} (\psi \vee \mathbf{G} \varphi) \quad (\text{Lemma 11.2}) \\ \iff w &\models \varphi \mathbf{W} \psi \quad (\text{Equation (10)}) \quad \square \end{aligned}$$

Proposition 21 gives us all we need to define a formula $\varphi[M]_2^\Sigma$ satisfying Equation (8).

Definition 22. *Let φ be a formula and let $M \subseteq \mu(\varphi)$. The formula $\varphi[M]_2^\Sigma$ is inductively defined as follows for $\mathbf{R}$ and $\mathbf{W}$*

$$(\varphi \mathbf{R} \psi)[M]_2^\Sigma = (\varphi[M]_2^\Sigma \vee \mathbf{G}(\psi[M]_1^\Pi)) \mathbf{M} \psi[M]_2^\Sigma$$

$$(\varphi \mathbf{W} \psi)[M]_2^\Sigma = \varphi[M]_2^\Sigma \mathbf{U} (\psi[M]_2^\Sigma \vee \mathbf{G}(\varphi[M]_1^\Pi))$$

and homomorphically for all other cases.

A straightforward induction on φ shows that $\varphi[M]_2^\Sigma \in \Sigma_2$, justifying our notation. We prove that $\varphi[M]_2^\Sigma$ satisfies (8) by checking that it satisfies the conditions of Lemma 17.

Theorem 23. *Let φ be a formula. Then:*

$$\varphi \equiv \bigvee_{\substack{M \subseteq \mu(\varphi) \\ N \subseteq v(\varphi)}} \left(\varphi[M]_2^\Sigma \wedge \bigwedge_{\psi \in M} \mathbf{GF}(\psi[N]_1^\Sigma) \wedge \bigwedge_{\psi \in N} \mathbf{FG}(\psi[M]_1^\Pi) \right)$$

Proof. We show that conditions (a) and (b) of Lemma 17 hold.

(a) The proof is an easy induction on φ , applying Lemma 18 where necessary.

(b) We prove that

$$\forall w. w \models \varphi \iff w \models \varphi[\mathcal{GF}_w^\varphi]_2^\Sigma \quad (12)$$

holds by structural induction on φ . We make use of the identity

$$\psi[M]_2^\Sigma = \psi[M \cap \mu(\psi)]_2^\Sigma \quad (13)$$

which follows immediately from the fact that formulas in $M \setminus \mu(\psi)$ are not subformulas of ψ .

The base of the induction is $\varphi \in \{\mathbf{tt}, \mathbf{ff}, a, \neg a\}$. In all these cases we have $\varphi = \varphi[\mathcal{GF}_w^\varphi]_2^\Sigma$ by definition, and so (12) holds vacuously. All other cases in which $\varphi[M]_2^\Sigma$ is defined homomorphically are handled in the same way. We consider only one of them:

Case $\varphi = \psi_1 \mathbf{U} \psi_2$. By assumption, the induction hypothesis (12) holds for ψ_1 and ψ_2 , giving:

$$\forall u. (u \models \psi_1 \iff u \models \psi_1[\mathcal{GF}_u^{\psi_1}]_2^\Sigma) \quad (14)$$

$$\forall v. (v \models \psi_2 \iff v \models \psi_2[\mathcal{GF}_v^{\psi_2}]_2^\Sigma) \quad (15)$$

In order to use these two equivalences for the induction step, we need to replace $\mathcal{GF}_u^{\psi_1}$ and $\mathcal{GF}_v^{\psi_2}$ by $\mathcal{GF}_w^\varphi$ in the context of $[\cdot]_2^\Sigma$. For this we instantiate $u := w_i$ and $v := w_j$ for arbitrary $i, j \geq 0$ in (14) and (15). With this choice u and v are suffixes of w , and so thus we get $\mathcal{GF}_u^{\psi_1} = \mathcal{GF}_v^{\psi_2} = \mathcal{GF}_w^\varphi$. Notice further that, by intersection with $\mu(\cdot)$, we have $\mathcal{GF}_u^{\psi_1} = \mathcal{GF}_w^\varphi \cap \mu(\psi_1)$ and $\mathcal{GF}_v^{\psi_2} = \mathcal{GF}_w^\varphi \cap \mu(\psi_2)$. From (13) we obtain:

$$\forall i. (w_i \models \psi_1 \iff w_i \models \psi_1[\mathcal{GF}_w^\varphi]_2^\Sigma) \quad (16)$$

$$\forall j. (w_j \models \psi_2 \iff w_j \models \psi_2[\mathcal{GF}_w^\varphi]_2^\Sigma) \quad (17)$$

Applying (16) and (17) we get:

$$\begin{aligned} w &\models \psi_1 \mathbf{U} \psi_2 \\ \iff \exists k. w_k &\models \psi_2 \wedge (\forall \ell < k. w_\ell \models \psi_1) \\ \iff \exists k. w_k &\models \psi_2[\mathcal{GF}_w^\varphi]_2^\Sigma \wedge (\forall \ell < k. w_\ell \models \psi_1[\mathcal{GF}_w^\varphi]_2^\Sigma) \\ \iff w &\models (\psi_1 \mathbf{U} \psi_2)[\mathcal{GF}_w^\varphi]_2^\Sigma \end{aligned}$$

which concludes the proof.

The remaining cases are $\varphi = \psi_1 \mathbf{R} \psi_2$ and $\varphi = \psi_1 \mathbf{W} \psi_2$. Again, we only consider one of them, the other one being analogous.

Case $\varphi = \psi_1 \mathbf{W} \psi_2$. The argumentation is only slightly more complicated than that of the $\psi_1 \mathbf{U} \psi_2$ case. By induction hypothesis (16) and (17) hold. With the help of Lemma 20 we derive:

$$\begin{aligned} w &\models \psi_1 \mathbf{W} \psi_2 \\ \iff w &\models \psi_1 \mathbf{U} (\psi_2 \vee \mathbf{G}(\psi_1[\mathcal{GF}_w^\varphi]_1^\Pi)) \quad (\text{Proposition 21}) \\ \iff w &\models \psi_1 \mathbf{U} (\psi_2 \vee \mathbf{G}(\psi_1[\mathcal{GF}_w^\varphi]_1^\Pi)) \\ &\quad (\psi[M]_1^\Pi = \psi[M \cap \mu(\psi)]_1^\Pi) \\ \iff w &\models \psi_1[\mathcal{GF}_w^\varphi]_2^\Sigma \mathbf{U} (\psi_2[\mathcal{GF}_w^\varphi]_2^\Sigma \vee \mathbf{G}(\psi_1[\mathcal{GF}_w^\varphi]_1^\Pi)) \\ &\quad ((16) \text{ and } (17)) \\ \iff w &\models (\psi_1 \mathbf{W} \psi_2)[\mathcal{GF}_w^\varphi]_2^\Sigma \end{aligned}$$

$\square$

Example 24. Let $\varphi = F(a \wedge G(b \vee Fc))$. We have $\mu(\varphi) = \{\varphi, Fc\}$ and $\nu(\varphi) = \{G(b \vee Fc)\}$, and so the right-hand-side of Theorem 23 has eight disjuncts. However, contrary to Example 16, we have $\varphi[M]_2^\Sigma \neq \mathbf{ff}$ for every $M \subseteq \{\varphi, Fc\}$. Let $\Phi(M, N)$ be the disjunct for given sets M, N . We consider two cases:

Case $M := \emptyset, N := \emptyset$. In this case $\Phi(\emptyset, \emptyset) = \varphi[\emptyset]_2^\Sigma$, because the conjunctions over M and N are vacuous. We have:

$$\begin{aligned} \Phi(\emptyset, \emptyset) &= \varphi[\emptyset]_2^\Sigma \\ &= F(a \wedge (G(b \vee Fc)[\emptyset]_2^\Sigma)) \\ &= F(a \wedge (((b \vee Fc) \mathbf{W} \mathbf{ff})[\emptyset]_2^\Sigma)) \\ &= F(a \wedge ((b \vee Fc)[\emptyset]_2^\Sigma \mathbf{U} (\mathbf{ff} \vee G((b \vee Fc)[\emptyset]_1^\Pi)))) \\ &= F(a \wedge ((b \vee Fc) \mathbf{U} Gb)) \end{aligned}$$

Case $M := \{Fc\}, N := \{G(b \vee Fc)\}$. We get:

$$\begin{aligned} \varphi[M]_2^\Sigma &= F(a \wedge ((b \vee Fc)[M]_2^\Sigma \mathbf{U} (\mathbf{ff} \vee G((b \vee Fc)[M]_1^\Pi)))) \\ &= F(a \wedge ((b \vee Fc) \mathbf{U} (\mathbf{ff} \vee \mathbf{tt}))) = Fa \end{aligned}$$

Further, we have $\mathbf{FG}(G(b \vee Fc)[M]_1^\Pi) = \mathbf{FG}(\mathbf{Gtt}) = \mathbf{tt}$ and $\mathbf{GF}((Fc)[N]_1^\Sigma) = \mathbf{GF}(Fc) = \mathbf{GFc}$. So in this case we obtain $\Phi(\{Fc\}, \{G(b \vee Fc)\}) = Fa \wedge \mathbf{GFc}$.

Repeating this process for all possible sets M, N and bringing the resulting formula in disjunctive normal form we finally get

$$\varphi \equiv F(a \wedge ((b \vee Fc) \mathbf{U} Gb)) \vee (Fa \wedge \mathbf{GFc})$$

6.1 Complexity of the Normalisation Procedure

We show that the normalisation procedure has at most single exponential blowup in the length of the formula, improving on the previously known non-elementary bound.

Proposition 25. Let φ be a formula with length n . Then there exists an equivalent formula φ_{Δ_2} in Δ_2 of length $2^{2n+O(1)}$.

Proof. Let ψ be an arbitrary formula. We let $|\psi|$ denote the length of formula and start by giving bounds on $\psi[M]_1^\Pi$, $\psi[N]_1^\Sigma$, and $\psi[M]_2^\Sigma$. For this let $M \subseteq \mu(\psi)$ and $N \subseteq \nu(\psi)$ be sets of formulas. We obtain by induction on the structure of ψ that $|\psi[M]_1^\Pi| \leq |\psi|$, $|\psi[N]_1^\Sigma| \leq |\psi|$, and $|\psi[M]_2^\Sigma| \leq 2^{|\psi|+1}$.

Consider now the right-hand side of Theorem 23 as the postulated φ_{Δ_2} . Using these bounds we calculate the maximal size of a disjunct and obtain:

$$2^{n+1} + n(n+3) + n(n+3) + 1 = 2^{n+1} + 2n^2 + 6n + 1$$

For sufficiently large n , i.e. $n > 5$, we can bound this by 2^{n+2} . There exist at most 2^n disjuncts and thus the formula is at most of size 2^{2n+2} for $n > 5$. $\square$

6.2 A Dual Normal Form

We obtained Theorem 23 by relying on the LTL equivalence (10) and (11) for $\mathbf{W}$ and $\mathbf{R}$. Using dual LTL-equivalences for

$\mathbf{U}$ and $\mathbf{M}$, $\varphi \mathbf{U} \psi \equiv (\varphi \wedge F\psi) \mathbf{W} \psi$ and $\varphi \mathbf{M} \psi \equiv \varphi \mathbf{R}(\psi \wedge F\varphi)$, we can also obtain a dual normalisation procedure:

Definition 26. Let φ be a formula and let $N \subseteq \nu(\varphi)$ be a set of formulas. The formula $\varphi[N]_2^\Pi$ is inductively defined as follows for $\mathbf{U}$ and $\mathbf{M}$:

$$\begin{aligned} (\varphi \mathbf{U} \psi)[N]_2^\Pi &= (\varphi[N]_2^\Pi \wedge F(\psi[N]_1^\Sigma)) \mathbf{W} \psi[N]_2^\Pi \\ (\varphi \mathbf{M} \psi)[N]_2^\Pi &= \varphi[N]_2^\Pi \mathbf{R}(\psi[N]_2^\Pi \wedge F(\varphi[N]_1^\Sigma)) \end{aligned}$$

and homomorphically for all other cases.

Theorem 27. Let φ be a formula. Then:

$$\varphi \equiv \bigvee_{\substack{M \subseteq \mu(\varphi) \\ N \subseteq \nu(\varphi)}} \left(\varphi[N]_2^\Pi \wedge \bigwedge_{\psi \in M} \mathbf{GF}(\psi[N]_1^\Sigma) \wedge \bigwedge_{\psi \in N} \mathbf{FG}(\psi[M]_1^\Pi) \right)$$

7 A Translation from LTL to Deterministic Rabin Automata (DRW)

We apply our Δ_2 -normalisation procedure to derive a new translation from LTL to DRW via weak alternating automata (AWW). While the previously existing normalisation procedures could also be used to translate LTL into DRW, the resulting DRW could have non-elementary size in the length of the formula, making them impractical. We show that, thanks to the single exponential blow-up of the new procedure, the new translation has double exponential blow-up, which is asymptotically optimal.

It is well-known [14, 24] that an LTL formula φ of length n can be translated into an AWW with $O(n)$ states. We show that, if φ is in normal form, i.e., a disjunction as in Theorem 23, then the AWW can be chosen so that every path through the automaton switches at most once between accepting and non-accepting states. We then prove that determinising AWWs satisfying this additional property is much simpler than the general case.

The section is structured as follows: Section 7.1 introduces basic definitions, Section 7.2 shows how to translate an Δ_2 -formula into AWWs with at most one switch, and Section 7.3 presents the determinisation procedure for this subclass of AWWs.

7.1 Weak and Very Weak Alternating Automata

Let X be a finite set. The set of positive Boolean formulas over X , denoted by $\mathcal{B}^+(X)$, is the closure of $X \cup \{\mathbf{tt}, \mathbf{ff}\}$ under disjunction and conjunction. A set $S \subseteq X$ is a model of $\theta \in \mathcal{B}^+(X)$ if the truth assignment that assigns true to the elements of S and false to the elements of $X \setminus S$ satisfies θ . Observe, that if S is a model of θ and $S \subseteq S'$ then S' is also a model. A model S is minimal if no proper subset of S is a model. The set of minimal models is denoted $\mathcal{M}_\theta$. Two formulas are equivalent, denoted $\theta \equiv \theta'$, if their set of minimal models is equal, i.e., $\mathcal{M}_\theta = \mathcal{M}_{\theta'}$.

Alternating automata. An alternating Büchi word automaton over an alphabet Σ is a tuple $\mathcal{A} = \langle \Sigma, Q, \theta_0, \delta, \alpha \rangle$, where Q is a finite set of states, $\theta_0 \in \mathcal{B}^+(Q)$ is an initial formula, $\delta: Q \times \Sigma \mapsto \mathcal{B}^+(Q)$ is the transition function, and $\alpha \subseteq Q$ is the acceptance condition. A *run* of $\mathcal{A}$ on the word w is a directed acyclic graph $G = (V, E)$ satisfying the following properties:

- $V \subseteq Q \times \mathbb{N}_0$, and $E \subseteq \bigcup_{l \geq 0} ((Q \times \{l\}) \times (Q \times \{l+1\}))$.
- There exists a minimal model S of θ_0 such that $(q, 0) \in V$ iff $q \in S$.
- For every $(q, l) \in V$, either $\delta(q, w[l]) \equiv \mathbf{ff}$ or the set $\{q' : ((q, l), (q', l+1)) \in E\}$ is a minimal model of $\delta(q, w[l])$.
- For every $(q, l) \in V \setminus (Q \times \{0\})$ there exists $q' \in Q$ such that $((q', l-1), (q, l)) \in E$.

Runs can be finite or infinite. A run G is *accepting* if

- (a) $\delta(q, w[l]) \not\equiv \mathbf{ff}$ for every $(q, l) \in V$, and
- (b) every infinite path of G visits α -nodes (that is, nodes (q, l) such that $q \in \alpha$) infinitely often.

In particular, every finite run satisfying (a) is accepting. $\mathcal{A}$ accepts a word w iff it has an accepting run G on w . The language $\mathcal{L}(\mathcal{A})$ recognised by $\mathcal{A}$ is the set of words accepted by $\mathcal{A}$. Two automata are equivalent if they recognise the same language.

Alternating *co-Büchi* automata are defined analogously, changing condition (b) by the co-Büchi condition (every infinite path of G only visits α -nodes finitely often). Finally, in alternating *Rabin* automata α is a set of *Rabin pairs* $(F, I) \subseteq Q \times Q$, and (b) is replaced by the Rabin condition (there exists a Rabin pair $(F, I) \in \alpha$ such that every infinite path visits states of F only finitely often and states of I infinitely often).

An automaton is *deterministic* if for every state $q \in Q$ and every letter $a \in \Sigma$ there exists $q' \in Q$ such that $\delta(q, a) = q'$, and *non-deterministic* if for every $q \in Q$ and every $a \in \Sigma$ there exists $Q' \subseteq Q$ such that $\delta(q, a) = \bigvee_{q' \in Q'} q'$.

The following definitions are useful for reasoning about runs: A set $U \subseteq Q$ is called a *level*. If $U \subseteq \alpha$, then U is an α -level. A level U' is a *successor* of U w.r.t. $a \in \Sigma$, also called a -successor, if for every $q \in U$ there is a minimal model S_q of $\delta(q, a)$ such that $U' = \bigcup_{q \in U} S_q$. The k -th level of a run $G = (V, E)$ is the set $\{q : (q, k) \in V\}$. Observe that a level can be empty, and empty levels are α -levels. Further, by definition a level has no successors w.r.t. a iff it contains a state q such that $\delta(q, a) \equiv \mathbf{ff}$. In particular, every level of an accepting run has at least one successor.

Weak and very weak automata. Let $\mathcal{A} = \langle \Sigma, Q, \theta_0, \delta, \alpha \rangle$ be an alternating (co-)Büchi automaton. We write $q \longrightarrow q'$ if there is $a \in \Sigma$ such that q' belongs to some minimal model of $\delta(q, a)$. $\mathcal{A}$ is *weak* if there is a partition $Q_1, \dots, Q_m$ of Q such that

- for every $q, q' \in Q$, if $q \longrightarrow q'$ then there are $i \leq j$ such that $q \in Q_i$ and $q' \in Q_j$, and

Figure 3. A1W for $\varphi = F(a \wedge XG(b \vee XF c))$ with $\Sigma = 2^{\{a,b,c\}}$, $\theta_0 = q_0$, and $\alpha = \{q_1\}$.

- for every $1 \leq i \leq m$: $Q_i \subseteq \alpha$ or $Q_i \cap \alpha = \emptyset$.

$\mathcal{A}$ is *very weak* or *linear* if it is weak and every class Q_i of the partition is a singleton ($|Q_i| = 1$). We let AWW and A1W denote the set of weak and very weak alternating automata, respectively. Observe that for every weak automaton with a co-Büchi acceptance condition we can define a Büchi acceptance condition on the same structure recognising the same language. Thus we will from now on assume that every weak automaton is equipped with a Büchi acceptance condition.

We define the height of a weak alternating automaton. The definition is very similar, but not identical, to the one of [6]. A weak automaton $\mathcal{A}$ has *height* n if every path $q \rightarrow q' \rightarrow q'' \dots$ of $\mathcal{A}$ alternates at most $n - 1$ times between α and $Q \setminus \alpha$. For example, the automaton in Figure 3 has height 3. We let $AWW[n]$ (A1W[n]) denote the sets of all (very-)weak alternating automata with height at most n . Further, we let $AWW[n, \mathbf{A}]$ (resp. $AWW[n, \mathbf{R}]$) denote the set of automata of $AWW[n]$ whose initial formula satisfies $\theta_0 \in \mathcal{B}(\alpha)^+$ (resp. $\theta_0 \in \mathcal{B}(Q \setminus \alpha)^+$). For example the automaton of Figure 3 belongs to $A1W[3, \mathbf{R}]$.

7.2 Translation of LTL to A1W[2]

In the standard translation [24] of LTL to A1W, the states of the A1W for a formula φ are subformulas of φ , or negations thereof. We show that, at the price of a slightly more complicated translation, the resulting A1W for a Δ_i -formula belongs to $A1W[i]$. Thus by using Theorem 23 every LTL formula can be translated to an $A1W[2]$. The idea of the construction is to use subformulas as states ensuring that

1. the transition relation can only lead from a formula to another formula at the same level or a lower level in the syntactic-future hierarchy (Figure 1b), and
2. accepting states are Π_i subformulas.

This immediately leads to “at most one alternation”. However, there is a little technical problem: the level of a formula is not always well-defined, because some formulas do not belong to one single lowest level of the hierarchy. For

example, Xa belongs to both Π_1 and Σ_1 . So we need a mechanism to *disambiguate* these states. Formally we proceed as follows:

A formula is *proper* if it is neither a Boolean constant (**tt**, **ff**) nor a conjunction or disjunction. A *state* in our modified translation is an expression of the form $\langle \psi \rangle_\Gamma$, where ψ is a proper formula, and Γ is a smallest class of the syntactic-future hierarchy without the zeroth-level (Definition 5) that contains ψ . Hence we start with the classes Σ_1 and Π_1 and Γ lies strictly above Δ_0 . Observe that for some formulas there is more than one smallest class. For example, since $Xa \in \Sigma_1 \cap \Pi_1$, both Σ_1 and Π_1 are smallest classes containing Xa , and so both $\langle Xa \rangle_{\Sigma_1}$ and $\langle Xa \rangle_{\Pi_1}$ are states. For other formulas the class is unique. For example, the only state for aWb is $\langle aWb \rangle_{\Pi_1}$.

We assign to every formula ψ of LTL and every class Γ a Boolean combination of states, denoted $[\psi]_{\leq \Gamma}$, as follows:

- $[\mathbf{tt}]_{\leq \Gamma} = \mathbf{tt}$ and $[\mathbf{ff}]_{\leq \Gamma} = \mathbf{ff}$.
- $[\psi_1 \vee \psi_2]_{\leq \Gamma} = [\psi_1]_{\leq \Gamma} \vee [\psi_2]_{\leq \Gamma}$
- $[\psi_1 \wedge \psi_2]_{\leq \Gamma} = [\psi_1]_{\leq \Gamma} \wedge [\psi_2]_{\leq \Gamma}$
- If ψ is a proper formula, then $[\psi]_{\leq \Gamma} = \bigvee_{\Gamma' \leq \Gamma} \langle \psi \rangle_{\Gamma'}$, where $\Gamma' \leq \Gamma$ means that $\Gamma' = \Gamma$ or Γ' is below Γ .

For example, we obtain $[Xa]_{\leq \Sigma_2} = \langle Xa \rangle_{\Sigma_1} \vee \langle Xa \rangle_{\Pi_1}$ and $[Xa]_{\leq \Sigma_1} = \langle Xa \rangle_{\Sigma_1}$. Moreover, $[Fa]_{\leq \Pi_1} = \mathbf{ff}$, since there is no $\Gamma' \leq \Pi_1$ such that $Fa \in \Gamma'$.

Let $\varphi \in \Delta_i$ for some $i \geq 0$, and let $\text{sf}(\varphi)$ be the set of proper subformulas of φ . The automaton $\mathcal{A}_\varphi = \langle 2^{A\varphi}, Q, \theta_0, \delta, \alpha \rangle$ is defined as follows:

- $Q = \{ \langle \psi \rangle_\Gamma : \psi \in \text{sf}(\varphi), \Gamma \leq \Delta_i \}$.
- $\theta_0 = [\varphi]_{\leq \Delta_i}$.
- $\alpha = \{ \langle \psi \rangle_{\Pi_i} \in Q : i > 0 \}$.
- δ is the restriction to $Q \times \Sigma$ of the function $\delta : \mathcal{B}^+(Q) \times \Sigma \rightarrow \mathcal{B}^+(Q)$ (notice that we overload δ) defined inductively as follows:

$$\begin{aligned} \delta(\langle a \rangle_\Gamma, \sigma) &= \text{if } a \in \sigma \text{ then } \mathbf{tt} \text{ else } \mathbf{ff} \\ \delta(\langle \neg a \rangle_\Gamma, \sigma) &= \text{if } a \notin \sigma \text{ then } \mathbf{tt} \text{ else } \mathbf{ff} \\ \delta(\langle X\psi \rangle_\Gamma, \sigma) &= [\psi]_{\leq \Gamma} \\ \delta(\langle \varphi U \psi \rangle_\Gamma, \sigma) &= \delta([\psi \vee (\varphi \wedge X(\varphi U \psi))]_{\leq \Gamma}, \sigma) \\ \delta(\langle \varphi W \psi \rangle_\Gamma, \sigma) &= \delta([\psi \vee (\varphi \wedge X(\varphi W \psi))]_{\leq \Gamma}, \sigma) \\ \delta(\langle \varphi R \psi \rangle_\Gamma, \sigma) &= \delta([\psi \wedge (\varphi \vee X(\varphi R \psi))]_{\leq \Gamma}, \sigma) \\ \delta(\langle \varphi M \psi \rangle_\Gamma, \sigma) &= \delta([\psi \wedge (\varphi \vee X(\varphi M \psi))]_{\leq \Gamma}, \sigma) \end{aligned}$$

All other cases (**tt**, **ff**, $\wedge$, and $\vee$) are defined homomorphically. Observe that the Γ -bound for the **U**, **W**, **R**, and **M** cases suffice, since every Γ is closed under conjunction, disjunction and application of **X**.

An example of this construction is displayed in Figure 3. The states are labelled $q_0 = \langle \varphi \rangle_{\Sigma_3}$, $q_1 = \langle G(b \vee XFc) \rangle_{\Pi_2}$, and $q_2 = \langle Fc \rangle_{\Sigma_1}$.

Lemma 28. *Let φ be a formula of Δ_i . The automaton $\mathcal{A}_\varphi$ belongs to $A1W[i]$, has $2|\text{sf}(\varphi)|$ states, and recognises $\mathcal{L}(\varphi)$.*

Proof. Let us first show that $\mathcal{A}_\varphi$ belongs to $A1W[i]$. It follows immediately from the definition of $\mathcal{A}_\varphi$ that for every two states $\langle \psi \rangle_\Gamma, \langle \psi' \rangle_{\Gamma'}$ of $\mathcal{A}_\varphi$, if $\langle \psi \rangle_\Gamma \rightarrow \langle \psi' \rangle_{\Gamma'}$ then $\Gamma' \leq \Gamma$. So in every path there are at most $(i - 1)$ alternations between Σ and Π classes. Since the states of α are those annotated with Π classes, there are also at most $(i - 1)$ alternations between α and non- α states in a path.

To show that $\mathcal{A}_\varphi$ has at most $2|\text{sf}(\varphi)|$ states, observe that for every formula ψ there are at most two smallest classes of the syntactic-future hierarchy containing ψ . So $\mathcal{A}_\varphi$ has at most two states for each formula of $\text{sf}(\varphi)$.

To prove that $\mathcal{A}_\varphi$ recognises $\mathcal{L}(\varphi)$ one shows by induction on ψ that $\mathcal{A}_\varphi$ recognises $\mathcal{L}(\psi)$ from every Boolean combination of states $[\psi]_{\leq \Gamma}$ such that $\psi \in \Gamma$. The proof is completely analogous to the one appearing in [24]. $\square$

7.3 Determinisation of AWW[2]

We present a determinisation procedure for $AWW[2, R]$ and $AWW[2, A]$ inspired by the break-point construction from [13]. We only describe the construction for $AWW[2, R]$, as the one for $AWW[2, A]$ is dual. The following lemma states the key property of $AWW[2, R]$:

Lemma 29. *Let $\mathcal{A}$ be an $AWW[2, R]$. $\mathcal{A}$ accepts a word w if and only if there exists a run $G = (V, E)$ of $\mathcal{A}$ on w such that*

- $\delta(q, w[l]) \neq \mathbf{ff}$ for every $(q, l) \in V$, and
- there is a threshold $k \geq 0$ such that for every $l \geq k$ and for every node $(q, l) \in V$ the state q is accepting.

Proof. Assume that $\mathcal{A}$ accepts w . Let $G = (V, E)$ be an accepting run of $\mathcal{A}$ on w . Since $\mathcal{A}$ is an $AWW[2, R]$, every path has by definition at most *one* alternation of accepting and rejecting states and all states occurring in the initial formula are marked as rejecting. Hence if a node $(q, l) \in V$ is accepting, i.e. $q \in \alpha$, then all its descendants are accepting. Let $V_r \subseteq V$ be the set of rejecting nodes of V , i.e., the nodes $(q, l) \in V$ such that $q \notin \alpha$. Since the descendants of accepting nodes are accepting, the subgraph $G_r = (V_r, E \cap (V_r \times V_r))$ is acyclic and connected. If V_r is infinite, then by König's lemma G_r has an infinite path of non-accepting nodes, contradicting that G is an accepting run. So G_r is finite, and we can choose the threshold k as the largest level of a node of V_r , plus one.

Assume such a run $G = (V, E)$ exists. Condition (a) of an accepting run holds by hypothesis. For condition (b), just observe that, since the descendants of accepting nodes are accepting, and every infinite path of G contains a node of the form (q, k) , where k is the threshold level, every infinite path visits accepting nodes infinitely often. $\square$

However, Lemma 29 does not hold for $AWW[3, R]$:

Example 30. *Let $\mathcal{A}$ be the automaton shown in Figure 3 and let $w = \{a\}(\{b\}\{c\})^\omega$. Observe that $\mathcal{A}$ accepts w . We prove by contradiction that no run of $\mathcal{A}$ on w satisfies the properties described in Lemma 29. Assume such a run exists. By the*

definition of δ , the run must be infinite. Further, by assumption there exists a threshold k such that all successor levels of the run are exactly $\{q_1\}$. But there exists $k' > k$ such that $w[k'] = \{c\}$. Since $\delta(q_1, \{c\}) = q_1 \wedge q_2$, the $(k' + 1)$ -th level of the run contains q_2 . Contradiction.

Given an automaton $\mathcal{A}$ from AWW[2, R], we construct a deterministic co-Büchi automaton $\mathcal{D}$ such that $L(\mathcal{A}) = L(\mathcal{D})$. A state of the DCW $\mathcal{D}$ is a pair $(Levels, Promising)$, where $Levels \subseteq 2^Q$ and $Promising \subseteq 2^\alpha \cap Levels$. It follows that $\mathcal{D}$ has at most 3^{2^n} states. Intuitively, after reading a finite word $w_{0k} = a_0 \dots a_k$ the automaton $\mathcal{D}$ is in the state $(Levels_k, Promising_k)$, where $Levels_k$ contains the k -th levels of every run of $\mathcal{A}$ on all words with w_{0k} as prefix, and $Promising_k \subseteq Levels_k$ contains the α -levels of $Levels_k$ that can still “generate” an accepting run. For this, when $\mathcal{D}$ reads a_{i+1} , it moves from $(Levels_i, Promising_i)$ to $(Levels_{i+1}, Promising_{i+1})$, where $Levels_{i+1}$ contains the successors w.r.t. a_{i+1} of $Levels_i$, and $Promising_{i+1}$ is defined as follows:

- If $Promising_i \neq \emptyset$, then $Promising_{i+1}$ contains the successors w.r.t. a_{i+1} of $Promising_i$.
- If $Promising_i = \emptyset$, then $Promising_{i+1}$ contains the α -levels of $Levels_{i+1}$.

Finally, the co-Büchi condition contains the states $(Levels, Promising)$ such that $Promising = \emptyset$.

Intuitively, during its run on a word w , the automaton $\mathcal{D}$ tracks the promising levels, removing those without successors, because they can no longer produce an accepting run. If the $Promising$ set becomes empty infinitely often, then every run of $\mathcal{A}$ on w contains a level without successors, and so $\mathcal{A}$ does not accept w . If after some number of steps, say k , the $Promising$ set never becomes empty again, then $\mathcal{A}$ has a run on w such that every level is an α -level and has at least one successor, and so this run is accepting.

For the formal definition of $\mathcal{D}$ it is convenient to identify subsets of 2^Q and 2^α with formulas of $\mathcal{B}^+(Q)$, $\mathcal{B}^+(\alpha)$ (i.e., we identify a formula and its set of models). Further, we lift $\delta: Q \times \Sigma \mapsto \mathcal{B}(Q)^+$ to $\delta: \mathcal{B}^+(Q) \times \Sigma \mapsto \mathcal{B}^+(Q)$ in the canonical way. Finally, given $\varphi \in \mathcal{B}^+(Q)$ and $S \subseteq Q$, we let $\varphi[\mathbf{ff}/S]$ denote the result of substituting $\mathbf{ff}$ for every state of $Q \setminus \alpha$ in $\delta(q, a)$. With these notations, the deterministic Büchi automaton $\mathcal{D}$ equivalent to $\mathcal{A}$ can be described in four lines: $\mathcal{D} = \langle \Sigma, Q', q'_0, \delta', \alpha' \rangle$, where $Q' = \mathcal{B}^+(Q) \times \mathcal{B}^+(\alpha)$, $q'_0 = (\theta_0, \mathbf{ff})$, $\alpha' = \{(\theta, \mathbf{ff}) : \theta \in \mathcal{B}^+(Q)\}$, and

$$\delta'((q, p), a) = \begin{cases} (\delta(q, a), \delta(p, a)) & \text{if } p \neq \mathbf{ff} \\ (\delta(q, a), \delta(q, a)[\mathbf{ff}/Q \setminus \alpha]) & \text{otherwise.} \end{cases}$$

Lemma 31. *For every $\mathcal{A} \in \text{AWW}[2, R]$ with n states, the deterministic co-Büchi automaton $\mathcal{D}$ defined above satisfies $L(\mathcal{A}) = L(\mathcal{D})$, and has 3^{2^n} states. Dually, for every $\mathcal{A}' \in \text{AWW}[2, A]$ with n' states, there exists a deterministic Büchi automaton $\mathcal{D}'$ that has $3^{2^{n'}}$ states and that satisfies $L(\mathcal{A}') = L(\mathcal{D}')$.*

Proof. Assume w is accepted by $\mathcal{A}$. Let $G = (V, E)$ be an accepting run of $\mathcal{A}$ on w . By Lemma 29 there exists an index k such that all levels of G after the k -th one are contained in α and have at least one successor. Therefore, the run $(Levels_0, Promising_0), (Levels_1, Promising_1) \dots$ of $\mathcal{D}$ on w satisfies $Promising_i \neq \emptyset$ for almost all i , and so $\mathcal{D}$ accepts.

Assume w is accepted by $\mathcal{D}$. Let $(Levels_0, Promising_0), (Levels_1, Promising_1) \dots$ be the run of $\mathcal{D}$ on w . By definition, there is a $k \geq 0$ such that $Promising_i \neq \emptyset$ for every $i \geq k$. Choose levels $U_0, U_1, \dots, U_k$ such that

- $U_k \in Promising_k$, and
- for every $1 \leq i \leq k$, choose U_{i-1} as a predecessor of U_i (this is always possible by the definition of δ').

Further, for every $i \geq k$ choose U_{i+1} as a successor of U_i . Now, let $G = (V, E)$ be the graph given by

- for every $l \geq 0$, $(q, l) \in V$ iff $q \in U_l$; and
- $((q, l), (q', l+1)) \in E$ iff $q \in U_l$ and $q' \in S_q$, where S_q is the minimal model of $\delta(q, w[l])$ used in the definition of successor level.

It follows immediately from the definitions that G is an accepting run of $\mathcal{A}$. The second part is proven by complementing $\mathcal{A}'$, applying the just described construction, and replacing the co-Büchi condition by a Büchi condition. $\square$

This result leads to a determinisation procedure for AWW[2].

Lemma 32. *For every $\mathcal{A} = \langle \Sigma, Q, \theta_0, \delta, \alpha \rangle \in \text{AWW}[2]$ with $n = |Q|$ states and $m = |\mathcal{M}_{\theta_0}|$ minimal models of θ_0 there exists an equivalent deterministic Rabin automaton $\mathcal{D}$ with $2^{2^{n+\log_2 m+2}}$ states and with m Rabin pairs.*

Proof. Let $\mathcal{A} = \langle \Sigma, Q, \theta_0, \delta, \alpha \rangle$. Given $Q' \subseteq Q$, let $\mathcal{A}_{Q'}$ be the AWW[2] obtained from $\mathcal{A}$ by substituting $\bigwedge_{q \in Q'} q$ for the initial formula θ_0 . We claim that for each minimal model $S \in \mathcal{M}_{\theta_0}$ we can construct a deterministic Rabin automaton (DRW) $\mathcal{D}_S$ with at most $2^{2^{n+2}}$ states and a single Rabin pair, recognising the same language as $\mathcal{A}_S$. Let us first see how to construct $\mathcal{D}$, assuming the claim holds. By the claim we have $\mathcal{L}(\mathcal{A}) = \bigcup_{S \in \mathcal{M}_{\theta_0}} \mathcal{L}(\mathcal{A}_S)$. So we define $\mathcal{D}$ as the union of all the automata $\mathcal{D}_S$. Recall that given two DRWs with n_1, n_2 states and p_1, p_2 Rabin pairs we can construct a DRW for the union of their languages with $n_1 \times n_2$ states and $n_1 + n_2$ pairs. Since θ_0 has m models, $\mathcal{D}$ has at most m Rabin pairs and $(2^{2^{n+2}})^m = 2^{2^{n+\log_2 m+2}}$ states.

It remains to prove the claim. Partition S into $S \cap \alpha$ and $S \setminus \alpha$. We have $\mathcal{A}_{S \cap \alpha} \in \text{AWW}[2, A]$ and $\mathcal{A}_{S \setminus \alpha} \in \text{AWW}[2, R]$. By Lemma 31 there exists a deterministic Büchi automaton $\mathcal{D}_{S \cap \alpha}$ and a deterministic co-Büchi automaton $\mathcal{D}_{S \setminus \alpha}$ equivalent to $\mathcal{A}_{S \cap \alpha}$ and $\mathcal{A}_{S \setminus \alpha}$, respectively, both with at most 3^{2^n} states. Intersecting these two automata yields a deterministic Rabin automaton with at most $3^{2^{n+1}} \leq 2^{2^{n+2}}$ states and a single Rabin pair, and we are done. $\square$

7.4 Translation of LTL to DRW

We combine the normalisation procedure and the translation of LTL to A1W of the previous section to obtain for every formula of LTL an equivalent DRW of double exponential size. Given a formula φ we have: $\varphi \equiv \bigvee_{\substack{M \subseteq \mu(\varphi) \\ N \subseteq \nu(\varphi)}} \varphi_{M,N}$

where

$$\varphi_{M,N} = \left(\varphi[M]_2^\Pi \wedge \bigwedge_{\psi \in M} \mathbf{GF}(\psi[N]_1^\Sigma) \wedge \bigwedge_{\psi \in N} \mathbf{FG}(\psi[M]_1^\Pi) \right)$$

Using the results of Section 7.2, we translate each formula $\varphi_{M,N}$ to an A1W[2], and then, applying the determinisation algorithm of Section 7.3, to a DRW. Finally, using the well-known union operation for DRWs, we obtain a DRW for φ .

In order to bound the number of states of the final DRW, we first need to determine the number of states of the A1W for each $\varphi_{M,N}$.

Lemma 33. *Let φ be a formula. For every $M \subseteq \mu(\varphi)$ and $N \subseteq \nu(\varphi)$, there exists an A1W[2] with $O(|sf(\varphi)|)$ states that recognises $\mathcal{L}(\varphi_{M,N})$.*

Proof. By Lemma 28, some A1W[2] with $O(|sf(\varphi_{M,N})|)$ states recognises $\mathcal{L}(\varphi_{M,N})$. So it suffices to show that $|sf(\varphi_{M,N})| \in O(|sf(\varphi)|)$, which follows from these claims, proved in the appendix of the extended version of this paper [23]:

1. $|\bigcup \{sf(\psi[M]_1^\Pi) : \psi \in sf(\varphi)\}| \leq |sf(\varphi)|$;
2. $|\bigcup \{sf(\psi[N]_1^\Sigma) : \psi \in sf(\varphi)\}| \leq |sf(\varphi)|$;
3. $|sf(\varphi[M]_2^\Pi)| \leq 3|sf(\varphi)|$. □

Proposition 34. *Let φ be a formula with n proper subformulas. There exists a deterministic Rabin automaton recognising $\mathcal{L}(\varphi)$ with $2^{2^{O(n)}}$ states and 2^n Rabin pairs.*

Proof. By Lemma 33 the set $sf(\varphi_{M,N})$ has at most $O(n)$ elements for every M, N . Further, due to Lemma 28 the automaton $\mathcal{A}_{\varphi_{M,N}}$ belongs to A1W[2] and has at most $O(n)$ states. Applying the construction of Lemma 32 we obtain a DRW with $2^{2^{O(n)}}$ states and a single Rabin pair. Using the union operation for DRWs we obtain a DRW for φ with $\left(2^{2^{O(n)}}\right)^{2^n} = 2^{2^{O(n)}}$ states. □

Remark 35. *The construction of Lemma 31 is close to Miyano and Hayashi's translation of alternating automata to non-deterministic automata [13], and to Schneider's translation of Σ_2 formulas to deterministic co-Büchi automata [18, p.219], all based on the break-point idea.*

7.5 Determinisation of Lower Classes

We now determinise AWW[1]. A deterministic automaton is *terminal-accepting* if all states are rejecting except a single accepting sink with a self-loop, and *terminal-rejecting* if all states are accepting except a single rejecting sink with a self-loop. It is easy to see that terminal-accepting and terminal-rejecting deterministic automata are closed under union and

intersection. When applied to AWW[1, A], the construction of Lemma 31, yields automata whose states have a trivial *Promising* set (either the empty set or the complete level). Further, the successor of an α -level is also an α -level. From these observations we easily get:

Corollary 36. *Let $\mathcal{A}$ be an automaton with n states.*

- *If $\mathcal{A} \in \text{AWW}[1, \mathbf{R}]$ (resp. $\mathcal{A} \in \text{AWW}[1, \mathbf{A}]$), then there exists a deterministic terminal-accepting (resp. terminal-rejecting) automaton recognising $\mathcal{L}(\mathcal{A})$ with 2^{2^n} states.*
- *If $\mathcal{A} \in \text{AWW}[1]$, then there exists deterministic weak automaton recognising $\mathcal{L}(\mathcal{A})$ with $2^{2^{n+\log_2 |\mathcal{M}_{\theta_0}|+1}}$ states.*

7.6 Preliminary Experimental Evaluation

We expect the LTL-to-DRW translation of this paper to produce automata similar in size (number of states, Rabin pairs) to the translations presented in [5, 21], which have been implemented using Owl [7] and have been extensively tested. Indeed, the “Master Theorem” of [5, 21] characterises the words satisfying a formula φ as those for which there exist sets M, N of subformulas satisfying three conditions, and so it has the same rough structure as our normal form. Further, for each disjunct of our normal form the automata constructions used in [5, 21] and the ones used in this paper are similar. Finally, in preliminary experiments we have compared the LTL-to-DRW translations from [21] and a prototype implementation, without optimisations, of the normalisation procedure of this paper. As benchmark sets we used the “Dwyer”-patterns [4], pre-processed as described in [21, Ch. 8], and the “Parametrised” formula set from [21, Ch. 8]. We observed that on the first set for 60% of the formulas the number of states of the resulting DRWs was equal, for 17% the number of states obtained using the construction of this paper was smaller, and for 23% the number of states was larger. On the second set the ratios were: 76% equal, 21% smaller, and 3% larger. For both sets combined we observed that in 85% of all 164 cases the difference in number of states was less than or equal to three.

We concluded that the main advantage of our translation is not its performance, but its modularity (it splits the procedure into a normalisation and a simplified translation phase) and its suitability for symbolic automata constructions. We leave a detailed experimental comparison and possible integration in Owl [7] (which in particular requires to examine different options for formula and automata simplification, as well as an extensive comparison to existing translations) for future work.

8 A Hierarchy of Alternating Weak and Very Weak Automata

The expressive power of weak and very weak alternating automata has been studied by Gurumurthy *et al.* in [6] and

Figure 4. Expressive power of AWWs after Gurumurthy *et al.* [6], and of A1Ws after Pelánek and Strejcek [16].

by Pelánek and Strejcek in [16], respectively. Both papers identify the number of alternations between accepting and non-accepting states as an important parameter, and define a hierarchy of automata classes based on it. Let $\text{AWW}_G[k]$ denote the class of AWW with at most $(k - 1)$ alternations defined in [6]. Similarly, let $\text{A1W}_{\text{PS}}[k, \mathbf{A}]$ and $\text{A1W}_{\text{PS}}[k, \mathbf{R}]$ denote the classes of A1W with at most $(k - 1)$ alternations and accepting or non-accepting initial state, respectively, defined in [16]. Finally, define $\text{A1W}_{\text{PS}}[k] = \text{A1W}_{\text{PS}}[k, \mathbf{A}] \cup \text{A1W}_{\text{PS}}[k, \mathbf{R}]$ ⁴. Figure 4 shows the results of [6] and [16]. We abuse language, and, for example, write $\Pi_2 = \text{A1W}_{\text{PS}}[2, \mathbf{A}]$ to denote that the class of languages satisfying formulas in Π_2 and the class of languages recognised by automata in $\text{A1W}_{\text{PS}}[2, \mathbf{A}]$ coincide.

Unfortunately, the results of [6] and [16] do not “match”. Due to slight differences in the definitions of height, e.g. the treatment of $\delta(\cdot) = \mathbf{ff}$ and $\delta(\cdot) = \mathbf{tt}$, the restriction to very weak automata of $\text{AWW}_G[k]$ does not match any class $\text{A1W}_{\text{PS}}[k']$ (that is, $\text{AWW}_G[k] \cap \text{A1W} \neq \text{A1W}_{\text{PS}}[k']$) and, vice versa, extending $\text{A1W}_{\text{PS}}[k]$ does not yield any $\text{AWW}_G[k']$. We show that our new definition of height unifies the two hierarchies, yielding the pleasant result shown in Figure 5. The result follows from Lemmas 28, 31 and 32, Corollary 36, and from constructions appearing in [6, 9, 16].

Proposition 37. $\text{AWW}[2] = \omega\text{-regular}$, $\text{AWW}[2, \mathbf{A}] = \text{DBW}$, $\text{AWW}[2, \mathbf{R}] = \text{DCW}$, $\text{AWW}[1] = \text{DWW}$, $\text{AWW}[1, \mathbf{A}] = \text{safety}$, $\text{AWW}[1, \mathbf{R}] = \text{co-safety}$, $\text{A1W}[1, \mathbf{R}] = \Sigma_1$, $\text{A1W}[1, \mathbf{A}] = \Pi_1$, $\text{A1W}[1] = \Delta_1$, $\text{A1W}[2, \mathbf{R}] = \Sigma_2$, $\text{A1W}[2, \mathbf{A}] = \Pi_2$, $\text{A1W}[2] = \Delta_2$.

Moreover, our single exponential normalisation procedure for LTL transfers to a single exponential normalisation procedure for A1W:

Lemma 38. *Let $\mathcal{A}$ be an A1W with n states over an alphabet with m letters. There exists $\mathcal{A}' \in \text{A1W}[2]$ with $2^{O(nm)}$ states such that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$.*

⁴In [16] the classes have different names.

Figure 5. Expressive power of AWWs and A1Ws

Proof. The translation from A1W to LTL used in Proposition 37 (an adaption of [9]) yields a formula $\chi_{\mathcal{A}}$ with at most $O(mn)$ proper subformulas. Applying our normalisation procedure to $\chi_{\mathcal{A}}$ yields an equivalent formula in Δ_2 with at most $2^{O(mn)}$ proper subformulas (Lemma 33). Applying Lemma 28 we obtain the postulated automaton $\mathcal{A}'$. $\square$

9 Conclusion

We have presented a purely syntactic normalisation procedure for LTL that transforms a given formula into an equivalent formula in Δ_2 , i.e., a formula with at most one alternation between least- and greatest-fixpoint operators. The procedure has single exponential blow-up, improving on the prohibitive non-elementary cost of previous constructions. The much better complexity of the new procedure (recall that normalisation procedures for CNF and DNF are also exponential) makes it attractive for its implementation and use in tools. We have presented a first promising application, namely a novel translation from LTL to DRW with double exponential blow-up. Finally, we have shown that the normalisation procedure for LTL can be transferred to a normalisation procedure for very weak alternating automata.

Currently we do not know if our normalisation procedure is asymptotically optimal. We conjecture that this is the case. For the translation of AWW to $\text{AWW}[2]$ we also have no further insight, besides the straightforward double exponential upper bound.

Acknowledgments

The authors want to thank Orna Kupferman for the suggestion to examine the expressive power of weak alternating automata and the anonymous reviewers for their helpful comments and remarks.

This work is partly funded by the German Research Foundation (DFG) project “Verified Model Checkers” (317422601) and partly funded by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme under grant agreement PaVeS (No 787-367).

References

- [1] Julian Brunner, Benedikt Seidl, and Salomon Sickert. 2019. A Verified and Compositional Translation of LTL to Deterministic Rabin Automata. In *10th International Conference on Interactive Theorem Proving, ITP 2019, September 9–12, 2019, Portland, OR, USA (LIPIcs, Vol. 141)*, John Harrison, John O’Leary, and Andrew Tolmach (Eds.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 11:1–11:19. <https://doi.org/10.4230/LIPIcs.ITP.2019.11>
- [2] Ivana Černá and Radek Pelánek. 2003. Relating Hierarchy of Temporal Properties to Model Checking. In *Mathematical Foundations of Computer Science 2003, 28th International Symposium, MFCS 2003, Bratislava, Slovakia, August 25–29, 2003, Proceedings (Lecture Notes in Computer Science, Vol. 2747)*, Branislav Rován and Peter Vojtáš (Eds.). Springer, 318–327. https://doi.org/10.1007/978-3-540-45138-9_26
- [3] Edward Y. Chang, Zohar Manna, and Amir Pnueli. 1992. Characterization of Temporal Property Classes. In *Automata, Languages and Programming, 19th International Colloquium, ICALP92, Vienna, Austria, July 13–17, 1992, Proceedings (Lecture Notes in Computer Science, Vol. 623)*, Werner Kuich (Ed.). Springer, 474–486. https://doi.org/10.1007/3-540-55719-9_97
- [4] Matthew B. Dwyer, George S. Avrunin, and James C. Corbett. 1998. Property specification patterns for finite-state verification. In *FMSP*. 7–15. <https://doi.org/10.1145/298595.298598>
- [5] Javier Esparza, Jan Kretínský, and Salomon Sickert. 2018. One Theorem to Rule Them All: A Unified Translation of LTL into ω -Automata. In *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09–12, 2018*, Anuj Dawar and Erich Grädel (Eds.). ACM, 384–393. <https://doi.org/10.1145/3209108.3209161>
- [6] Sankar Gurumurthy, Orna Kupferman, Fabio Somenzi, and Moshe Y. Vardi. 2003. On Complementing Nondeterministic Büchi Automata. In *Correct Hardware Design and Verification Methods, 12th IFIP WG 10.5 Advanced Research Working Conference, CHARME 2003, L’Aquila, Italy, October 21–24, 2003, Proceedings (Lecture Notes in Computer Science, Vol. 2860)*, Daniel Geist and Enrico Tronci (Eds.). Springer, 96–110. https://doi.org/10.1007/978-3-540-39724-3_10
- [7] Jan Kretínský, Tobias Meggendorfer, and Salomon Sickert. 2018. Owl: A Library for ω -Words, Automata, and LTL. In *Automated Technology for Verification and Analysis - 16th International Symposium, ATVA 2018, Los Angeles, CA, USA, October 7–10, 2018, Proceedings (Lecture Notes in Computer Science, Vol. 11138)*, Shuvendu K. Lahiri and Chao Wang (Eds.). Springer, 543–550. https://doi.org/10.1007/978-3-030-01090-4_34
- [8] Orna Lichtenstein, Amir Pnueli, and Lenore D. Zuck. 1985. The Glory of the Past. In *Logics of Programs, Conference, Brooklyn College, New York, NY, USA, June 17–19, 1985, Proceedings (Lecture Notes in Computer Science, Vol. 193)*, Rohit Parikh (Ed.). Springer, 196–218. https://doi.org/10.1007/3-540-15648-8_16
- [9] Christof Löding and Wolfgang Thomas. 2000. Alternating Automata and Logics over Infinite Words. In *Theoretical Computer Science, Exploring New Frontiers of Theoretical Informatics, International Conference IFIP TCS 2000, Sendai, Japan, August 17–19, 2000, Proceedings (Lecture Notes in Computer Science, Vol. 1872)*, Jan van Leeuwen, Osamu Watanabe, Masami Hagiya, Peter D. Mosses, and Takayasu Ito (Eds.). Springer, 521–535. https://doi.org/10.1007/3-540-44929-9_36
- [10] Oded Maler and Amir Pnueli. 1990. Tight Bounds on the Complexity of Cascaded Decomposition of Automata. In *31st Annual Symposium on Foundations of Computer Science, St. Louis, Missouri, USA, October 22–24, 1990, Volume II*. IEEE Computer Society, 672–682. <https://doi.org/10.1109/FSCS.1990.89589>
- [11] Zohar Manna and Amir Pnueli. 1990. A Hierarchy of Temporal Properties. In *Proceedings of the Ninth Annual ACM Symposium on Principles of Distributed Computing, Quebec City, Quebec, Canada, August 22–24, 1990*, Cynthia Dwork (Ed.). ACM, 377–410. <https://doi.org/10.1145/93385.93442>
- [12] Zohar Manna and Amir Pnueli. 1992. *The temporal logic of reactive and concurrent systems - specification*. Springer.
- [13] Satoru Miyano and Takeshi Hayashi. 1984. Alternating Finite Automata on omega-Words. *Theor. Comput. Sci.* 32 (1984), 321–330. [https://doi.org/10.1016/0304-3975\(84\)90049-5](https://doi.org/10.1016/0304-3975(84)90049-5)
- [14] David E. Muller, Ahmed Saoudi, and Paul E. Schupp. 1988. Weak Alternating Automata Give a Simple Explanation of Why Most Temporal and Dynamic Logics are Decidable in Exponential Time. In *Proceedings of the Third Annual Symposium on Logic in Computer Science (LICS ’88), Edinburgh, Scotland, UK, July 5–8, 1988*. IEEE Computer Society, 422–427. <https://doi.org/10.1109/LICS.1988.5139>
- [15] Tobias Nipkow, Lawrence C. Paulson, and Markus Wenzel. 2002. *Isabelle/HOL - A Proof Assistant for Higher-Order Logic*. Lecture Notes in Computer Science, Vol. 2283. Springer. <https://doi.org/10.1007/3-540-45949-9>
- [16] Radek Pelánek and Jan Strejcek. 2005. Deeper Connections Between LTL and Alternating Automata. In *Implementation and Application of Automata, 10th International Conference, CIAA 2005, Sophia Antipolis, France, June 27–29, 2005, Revised Selected Papers (Lecture Notes in Computer Science, Vol. 3845)*, Jacques Farré, Igor Litovsky, and Sylvain Schmitz (Eds.). Springer, 238–249. https://doi.org/10.1007/11605157_20
- [17] Nir Piterman and Amir Pnueli. 2018. Temporal Logic and Fair Discrete Systems. In *Handbook of Model Checking*, Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, and Roderick Bloem (Eds.). Springer, 27–73. https://doi.org/10.1007/978-3-319-10575-8_2
- [18] Klaus Schneider. 2004. *Verification of Reactive Systems - Formal Methods and Algorithms*. Springer. <https://doi.org/10.1007/978-3-662-10778-2>
- [19] Benedikt Seidl and Salomon Sickert. 2019. A Compositional and Unified Translation of LTL into ω -Automata. *Archive of Formal Proofs* 2019 (2019). https://www.isa-afp.org/entries/LTL_Master_Theorem.html
- [20] Salomon Sickert. 2016. Linear Temporal Logic. *Archive of Formal Proofs* 2016 (2016). <https://www.isa-afp.org/entries/LTL.shtml>
- [21] Salomon Sickert. 2019. *A Unified Translation of Linear Temporal Logic to ω -Automata*. Ph.D. Dissertation. Technical University of Munich, Germany. <http://nbn-resolving.de/urn:nbn:de:bvb:91-diss-20190801-1484932-1-4>
- [22] Salomon Sickert. 2020. An Efficient Normalisation Procedure for Linear Temporal Logic: Isabelle/HOL Formalisation. *Archive of Formal Proofs* 2020 (2020). https://www.isa-afp.org/entries/LTL_Normal_Form.html
- [23] Salomon Sickert and Javier Esparza. 2020. An Efficient Normalisation Procedure for Linear Temporal Logic and Very Weak Alternating Automata. *CoRR* abs/2005.00472 (2020). <http://arxiv.org/abs/2005.00472>
- [24] Moshe Y. Vardi. 1994. Nontraditional Applications of Automata Theory. In *Theoretical Aspects of Computer Software, International Conference TACS ’94, Sendai, Japan, April 19–22, 1994, Proceedings (Lecture Notes in Computer Science, Vol. 789)*, Masami Hagiya and John C. Mitchell (Eds.). Springer, 575–597. https://doi.org/10.1007/3-540-57887-0_116
- [25] Lenore D. Zuck. 1986. *Past Temporal Logic*. Ph.D. Dissertation. The Weizmann Institute of Science, Israel.