

LTL to Büchi Automata Translation: Fast and More Deterministic*

Tomáš Babiak, Mojmir Křetínský, Vojtěch Řehák, and Jan Strejček

Faculty of Informatics, Masaryk University
Botanická 68a, 60200 Brno, Czech Republic
{xbabiak, kretinsky, rehak, strejcek}@fi.muni.cz

Abstract. We introduce improvements in the algorithm by Gastin and Oddoux translating LTL formulae into Büchi automata via very weak alternating co-Büchi automata and generalized Büchi automata. Several improvements are based on specific properties of any formula where each branch of its syntax tree contains at least one *eventually* operator and at least one *always* operator. These changes usually result in faster translations and smaller automata. Other improvements reduce non-determinism in the produced automata. In fact, we modified all the steps of the original algorithm and its implementation known as LTL2BA. Experimental results show that our modifications are real improvements. Their implementations within an LTL2BA translation made LTL2BA very competitive with the current version of SPOT, sometimes outperforming it substantially.

1 Introduction

A translation of LTL formulae into equivalent Büchi automata plays an important role in many algorithms for LTL model checking, LTL satisfiability checking etc. For a long time, researchers aimed to find fast translations producing Büchi automata with a small number of states. This goal has led to the developments of several translation algorithms and many heuristics and optimizations including input formula reductions and optimizations of produced Büchi automata, see e.g. [3,4,9,18,11,12,10,7].

As the time goes, the translation objectives and their importance are changing. In particular, [17] demonstrates that for higher performance of the subsequent steps of the model checking process, it is more important to minimize the number of states with nondeterministic choice than the number of all states in resulting automata. Note that there are LTL formulae, e.g. FGa , for which no equivalent deterministic Büchi automaton exists. Further, model checking practice shows that one LTL formula is usually used in many different model checking tasks. Hence, it pays to invest enough computation time to get high quality (more

* The authors are supported by The Czech Science Foundation, grants 102/09/H042 (Babiak), 201/09/1389 (Křetínský), P202/10/1469 (Řehák, Strejček), P202/12/G061 (Křetínský, Řehák, Strejček), and P202/12/P612 (Řehák).

deterministic and/or minimal) automata as it may reduce computation time of many model checking tasks.

The new objectives lead to the developments of algorithms focusing on quality of produced automata. For example, [5] presents an effective algorithm translating LTL formulae of the fragment called *obligation* (see [14]) into *weak deterministic Büchi automata (WDBA)*. Moreover, WDBA can be minimized by the algorithm of [13]. There is also a SAT-based algorithm searching for minimal (nondeterministic) Büchi automata [8]. The main disadvantage of all the mentioned determinization and minimization algorithms is their long running time which limits their use.

Our research returns to the roots: we focus on a fast translation producing a relatively good output. This approach is justified by the following facts:

- The mentioned algorithms producing high quality automata often need, for a given LTL formula, some equivalent automaton as an input.
- The mentioned algorithms are usually feasible for short formulae only or for formulae with a simple structure.
- Given a fresh LTL formula, it can be useful to run vacuity checks, namely satisfiability of the formula and its negation, to detect bugs in the formula. In these checks, time of the LTL to automata translation can be more significant than time needed for subsequent computations (see [16]). Hence, we need a fast translator to support an early detection of bugs in formulae.

Considering the speed of an LTL to Büchi automata translation, LTL2BA [11] and SPOT [7] are two leading tools. Based on extensive experiments on LTL satisfiability checking, [16] even states:

The difference in performance between SPOT and LTL2BA, on one hand, and the rest of explicit tools is quite dramatic.

Each of the two tools is based on different algorithms.

In LTL2BA, the translation proceeds in three basic steps:

1. A given LTL formula is translated into a *very weak alternating automaton (VWAA)* with a co-Büchi accepting condition.
2. The alternating automaton is then translated into a *transition-based generalized Büchi automaton (TGBA)*, i.e. a generalized Büchi automaton with sets of accepting transitions instead of accepting states.
3. The generalized automaton is transformed (degeneralized) into a Büchi automaton (BA).

Each of the three automata is simplified during the translation.

SPOT translates a given LTL formula to a TGBA using a tableau method presented in [3]. The TGBA is then translated to a BA. Note that the model checking algorithm natively implemented in SPOT works directly with TGBAs. Prior to a translation, both LTL2BA and SPOT try to decrease the number of temporal operators in a given input formula by applications of reduction rules.

While the LTL to automata translation in SPOT is under the gradual development following the current trends (see [6] for improvements made in the last four years), LTL2BA underwent only one minor update in 2007 since its creation in 2001. In particular, SPOT reflects the changes in objectives. Therefore, SPOT usually produces more deterministic and smaller automata than LTL2BA, while LTL2BA is often a bit faster.

Our Contribution. We introduce several modifications of LTL2BA on both algorithmic and implementation levels. We suggest changes in all the steps of the translation algorithm. Our experimental results indicate that each modified step has a mostly positive effect on the translation. The new translator, called LTL3BA, is usually faster than the original LTL2BA and it produces smaller and more deterministic automata. Moreover, comparison of LTL3BA and the current version of SPOT (run without WDBA minimization that is very slow) shows that the produced automata are of similar quality and LTL3BA is usually faster.

Some modifications employ an observation that each LTL formula containing at least one *always* operator and at least one *eventually* operator on each branch of its syntax tree (with possible exceptions of branches going to the left subformula of any *until* or *release* operator) is prefix invariant. We call them *alternating* formulae. Indeed, validity of each alternating formula on a given word u depends purely on a suffix of u . In other words, it is not affected by any finite prefix of u . We apply this observation to construct new rules for formula reductions. Further, the observation justifies some changes in constructions of VWAA and TGBA. Intuitively, a state of a VWAA corresponds to a subformula that has to be satisfied by the rest of an accepted word. If the corresponding subformula is an alternating formula, then the state can be temporarily suspended for finitely many steps of the automaton.

Other changes in a VWAA construction are designed to lower nondeterminism. This is also a motivation for new simplification rules applied on intermediate automata. These rules remove some transitions of the automaton and hence reduce the number of nondeterministic choices in produced automata. The original simplification rules can be seen as special cases of the new rules. An effective implementation of this simplification required to change representation of transitions. Further, we add one ad-hoc modification speeding up the translation of selected (sub)formulae. Finally, we modify a simplification rule merging some states of resulting BA.

The rest of the paper is organized as follows. The next section recalls the definitions of LTL, VWAA, and TGBA, as presented in [11]. Section 3 focuses on alternating formulae and its properties. Sections 4, 5, 6, and 7 present new rules for formula reductions, modified translation of LTL to VWAA (including generalized simplification of VWAA), modified translation of VWAA to TGBA, and modified rule for simplification of BA, respectively. Finally, Section 8 is devoted to experimental results. The last section summarizes the achieved improvements.

2 Preliminaries

Linear Temporal Logic (LTL). The syntax of LTL [15] is defined as follows

$$\varphi ::= tt \mid a \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid X\varphi \mid \varphi U \varphi,$$

where tt stands for *true*, a ranges over a countable set AP of *atomic propositions*, X and U are temporal operators called *next* and *until*, respectively. The logic is interpreted over infinite words over the alphabet $\Sigma = 2^{AP'}$, where $AP' \subseteq AP$ is a finite subset. Given a word $u = u(0)u(1)u(2)\dots \in \Sigma^\omega$, by u_i we denote the i^{th} suffix of u , i.e. $u_i = u(i)u(i+1)\dots$.

The semantics of LTL formulae is defined inductively as follows:

$$\begin{aligned} u \models tt & \\ u \models a & \quad \text{iff } a \in u(0) \\ u \models \neg\varphi & \quad \text{iff } u \not\models \varphi \\ u \models \varphi_1 \vee \varphi_2 & \quad \text{iff } u \models \varphi_1 \text{ or } u \models \varphi_2 \\ u \models \varphi_1 \wedge \varphi_2 & \quad \text{iff } u \models \varphi_1 \text{ and } u \models \varphi_2 \\ u \models X\varphi & \quad \text{iff } u_1 \models \varphi \\ u \models \varphi_1 U \varphi_2 & \quad \text{iff } \exists i \geq 0. (u_i \models \varphi_2 \text{ and } \forall 0 \leq j < i. u_j \models \varphi_1) \end{aligned}$$

We say that a word u *satisfies* φ whenever $u \models \varphi$. Two formulae φ, ψ are *equivalent*, written $\varphi \equiv \psi$, if for each alphabet Σ and each $u \in \Sigma^\omega$ it holds $u \models \varphi \iff u \models \psi$. Given an alphabet Σ , a formula φ defines the language $L^\Sigma(\varphi) = \{u \in \Sigma^\omega \mid u \models \varphi\}$. We often write $L(\varphi)$ instead of $L^{2^{AP(\varphi)}}(\varphi)$, where $AP(\varphi)$ denotes the set of atomic propositions occurring in the formula φ .

We extend the LTL with derived temporal operators:

- $F\varphi$ called *eventually* and equivalent to $tt U \varphi$,
- $G\varphi$ called *always* and equivalent to $\neg F\neg\varphi$, and
- $\varphi R \psi$ called *release* and equivalent to $\neg(\neg\varphi U \neg\psi)$.

In the following, *temporal formula* is a formula where the topmost operator is neither conjunction, nor disjunction. A formula without any temporal operator is called *state formula*. Note that a and tt are both temporal and state formulae. An LTL formula is in *positive normal form* if no operator occurs in the scope of any negation. Each LTL formula can be easily transformed to positive normal form using De Morgan's laws for operators \vee and \wedge , equivalences for derived operators, and the following equivalences:

$$\neg(\varphi_1 U \varphi_2) \equiv \neg\varphi_1 R \neg\varphi_2 \quad \neg(\varphi_1 R \varphi_2) \equiv \neg\varphi_1 U \neg\varphi_2 \quad \neg X\varphi \equiv X\neg\varphi$$

Very Weak Alternating co-Büchi Automata (VWAA). A VWAA is a tuple $\mathcal{A} = (Q, \Sigma, \delta, I, F)$, where

- Q is a finite set of *states*, and we let $Q' = 2^Q$,
- Σ is a finite *alphabet*, and we let $\Sigma' = 2^\Sigma$,
- $\delta : Q \rightarrow 2^{\Sigma' \times Q'}$ is a *transition function*,

- $I \subseteq Q'$ is a set of *initial states*,
- $F \subseteq Q$ is a set of *accepting states*, and
- there exists a partial order on Q such that, for each state $q \in Q$, all the states occurring in $\delta(q)$ are lower or equal to q .

Note that the transition function δ uses Σ' instead of Σ . This enables to merge transitions that differ only by action labels. We sometimes use a propositional formula α over AP to describe the element $\{a \in \Sigma \mid a \text{ satisfies } \alpha\}$ of Σ' .

A *run* σ of VWAA \mathcal{A} over a word $w = w(0)w(1)w(2)\dots \in \Sigma^\omega$ is a labelled directed acyclic graph (V, E, λ) such that:

- V is partitioned into $\bigcup_{i=0}^{\infty} V_i$ with $E \subseteq \bigcup_{i=0}^{\infty} V_i \times V_{i+1}$,
- $\lambda : V \rightarrow Q$ is a labelling function,
- $\{\lambda(x) \mid x \in V_0\} \in I$, and
- for each $x \in V_i$, there exist $\alpha \in \Sigma'$, $q \in Q$ and $O \in Q'$ such that $w(i) \in \alpha$, $q = \lambda(x)$, $O = \{\lambda(y) \mid (x, y) \in E\}$, and $(\alpha, O) \in \delta(q)$.

A run σ is *accepting* if each branch in σ contains only finitely many nodes labelled by accepting states (co-Büchi acceptance condition). A word w is *accepted* if there is an accepting run over w .

Transition Based Generalized Büchi Automata (TGBA). A TGBA is a tuple $\mathcal{G} = (Q, \Sigma, \delta, I, \mathcal{F})$, where

- Q is a finite set of *states*,
- Σ is a finite *alphabet*, and we let $\Sigma' = 2^\Sigma$
- $\delta : Q \rightarrow 2^{\Sigma' \times Q}$ is a total *transition function*,
- $I \subseteq Q$ is a set of *initial states*, and
- $\mathcal{T} = \{T_1, T_2, \dots, T_m\}$ where $T_j \subseteq Q \times \Sigma' \times Q$ are sets of *accepting transitions*.

A run ρ of TGBA \mathcal{G} over a word $w = w(0)w(1)w(2)\dots \in \Sigma^\omega$ is a sequence of states $\rho = q_0q_1q_2\dots$, where $q_0 \in I$ is an initial state and, for each $i \geq 0$, there exists $\alpha \in \Sigma'$ such that $w(i) \in \alpha$ and $(\alpha, q_{i+1}) \in \delta(q_i)$. A run ρ is *accepting* if for each $1 \leq j \leq m$ it uses infinitely many transitions from T_j . A word w is *accepted* if there is an accepting run over w .

3 Alternating Formulae

We define the class of *alternating formulae* together with the classes of *pure eventuality* and *pure universality* formulae introduced in [9]. Let φ ranges over general LTL formulae. The classes of *pure eventuality* formulae μ , *pure universality* formulae ν , and *alternating* formulae ξ are defined as:

$$\begin{aligned}
\mu &::= F\varphi \mid \mu \vee \mu \mid \mu \wedge \mu \mid X\mu \mid \varphi U \mu \mid \mu R \mu \mid G\mu \\
\nu &::= G\varphi \mid \nu \vee \nu \mid \nu \wedge \nu \mid X\nu \mid \nu U \nu \mid \varphi R \nu \mid F\nu \\
\xi &::= G\mu \mid F\nu \mid \xi \vee \xi \mid \xi \wedge \xi \mid X\xi \mid \varphi U \xi \mid \varphi R \xi \mid F\xi \mid G\xi
\end{aligned}$$

Note that there are alternating formulae, e.g. $(a \text{ U } (\text{GF}b)) \wedge (c \text{ R } (\text{GF}d))$, that are neither pure eventuality formulae, nor pure universality formulae. Properties of the respective classes of formulae are summarized in the following lemmata.

Lemma 1. [9] *Let μ be a pure eventuality formula and ν be a pure universality formula. For all $u \in \Sigma^*$, $w \in \Sigma^\omega$ it holds:*

$$\begin{aligned} uw \models \mu &\iff w \models \mu \\ uw \models \nu &\implies w \models \nu \end{aligned}$$

Lemma 2. *Let ξ be an alternating formula. For all $u \in \Sigma^*$, $w \in \Sigma^\omega$ it holds:*

$$uw \models \xi \iff w \models \xi$$

In other words, pure eventuality formulae define *left-append closed* languages, pure universality formulae define *suffix closed* languages, and alternating formulae define *prefix-invariant* languages. The proof of Lemma 2 can be found in the full version of this paper [1].

Corollary 1. *Every alternating formula ξ satisfies $\xi \equiv X\xi$.*

Hence, in order to check whether w satisfies ξ it is possible to skip an arbitrary long finite prefix of the word w .

We use this property in new rule for formula reduction. Further, it has brought us to the notion of alternating formulae *suspension* during the translation of LTL to Büchi automata. We employ suspension on two different levels of the translation: the construction of a VWAA from an input LTL formula and the transformation of a VWAA into a TGBA.

4 Improvements in Reduction of LTL Formulae

Many rules reducing the number of temporal operators in an LTL formula have been presented in [18] and [9]. In this section we present some new reduction rules. For the rest of this section, φ, ψ range over LTL formulae and γ ranges over alternating ones.

$$\begin{array}{llll} X\varphi \text{ R } X\psi \equiv X(\varphi \text{ R } \psi) & \varphi \text{ U } \gamma \equiv \gamma & F\gamma \equiv \gamma & X\gamma \equiv \gamma \\ X\varphi \vee X\psi \equiv X(\varphi \vee \psi) & \varphi \text{ R } \gamma \equiv \gamma & G\gamma \equiv \gamma & \end{array}$$

The following equivalences are valid only on assumption that φ implies ψ .

$$\begin{array}{ll} \psi \text{ U } (\varphi \text{ U } \gamma) \equiv \psi \text{ U } \gamma & \varphi \wedge (\psi \wedge \gamma) \equiv (\varphi \wedge \gamma) \\ (\psi \text{ R } \gamma) \text{ R } \varphi \equiv \gamma \text{ R } \varphi & \psi \vee (\varphi \vee \gamma) \equiv (\psi \vee \gamma) \\ \varphi \text{ U } (\gamma \text{ R } (\psi \text{ U } \rho)) \equiv \gamma \text{ R } (\psi \text{ U } \rho) & \end{array}$$

Further, we have extended the set of rules deriving implications of the form $\varphi \Rightarrow \psi$. The upper formula is a precondition, the lower one is a conclusion.

$$\frac{G\varphi \Rightarrow \psi}{G\varphi \Rightarrow X\psi} \quad \frac{\varphi \Rightarrow F\psi}{X\varphi \Rightarrow F\psi} \quad \frac{\varphi \Rightarrow \psi}{X\varphi \Rightarrow X\psi}$$

5 Improvements in LTL to VWAA Translation

First, we recall the original translation of LTL to VWAA according to [11]. The translation utilizes two auxiliary operators:

- Let $\Sigma' = 2^\Sigma$, and let $Q' = 2^Q$. Given $J_1, J_2 \in 2^{\Sigma' \times Q'}$, we define

$$J_1 \otimes J_2 = \{(\alpha_1 \cap \alpha_2, O_1 \cup O_2) \mid (\alpha_1, O_1) \in J_1 \text{ and } (\alpha_2, O_2) \in J_2\}.$$

- Let $\overline{\psi}$ be an LTL formula in positive normal form. We define $\overline{\overline{\psi}}$ by:
 - $\overline{\psi} = \{\{\psi\}\}$ if ψ is a temporal formula,
 - $\overline{\psi_1 \wedge \psi_2} = \{\overline{O_1} \cup \overline{O_2} \mid O_1 \in \overline{\psi_1} \text{ and } O_2 \in \overline{\psi_2}\}$,
 - $\overline{\psi_1 \vee \psi_2} = \overline{\psi_1} \cup \overline{\psi_2}$.

Let φ be an LTL formula in positive normal form. An equivalent VWAA with a co-Büchi acceptance condition is constructed as $\mathcal{A}_\varphi = (Q, \Sigma, \delta, I, F)$, where Q is the set of temporal subformulae of φ , $\Sigma = 2^{AP(\varphi)}$, $I = \overline{\varphi}$, F is the set of all U-subformulae of φ , i.e formulae of the type $\psi_1 \text{ U } \psi_2$, and δ is defined as follows:

$$\begin{aligned} \delta(tt) &= \{(\Sigma, \emptyset)\} \\ \delta(p) &= \{(\Sigma_p, \emptyset)\} \text{ where } \Sigma_p = \{a \in \Sigma \mid p \in a\} \\ \delta(\neg p) &= \{(\Sigma_{\neg p}, \emptyset)\} \text{ where } \Sigma_{\neg p} = \Sigma \setminus \Sigma_p \\ \delta(\mathbf{X}\psi) &= \{(\Sigma, O) \mid O \in \overline{\psi}\} \\ \delta(\psi_1 \text{ U } \psi_2) &= \Delta(\psi_2) \cup (\Delta(\psi_1) \otimes \{(\Sigma, \{\psi_1 \text{ U } \psi_2\})\}) \\ \delta(\psi_1 \text{ R } \psi_2) &= \Delta(\psi_2) \otimes (\Delta(\psi_1) \cup \{(\Sigma, \{\psi_1 \text{ R } \psi_2\})\}) \\ \Delta(\psi) &= \delta(\psi) \text{ if } \psi \text{ is a temporal formula} \\ \Delta(\psi_1 \vee \psi_2) &= \Delta(\psi_1) \cup \Delta(\psi_2) \\ \Delta(\psi_1 \wedge \psi_2) &= \Delta(\psi_1) \otimes \Delta(\psi_2) \end{aligned}$$

Using the partial order “is a subformula of” on states of \mathcal{A}_φ , one can easily prove that \mathcal{A}_φ is very weak.

Improved Translation. In order to implement the suspension of alternating formulae, we modify the way the transition function δ handles the binary operators U, R, \vee , and \wedge . The original transition function δ reflects the following identities:

$$\varphi_1 \text{ U } \varphi_2 \equiv \varphi_2 \vee (\varphi_1 \wedge \mathbf{X}(\varphi_1 \text{ U } \varphi_2)) \quad \varphi_1 \text{ R } \varphi_2 \equiv \varphi_2 \wedge (\varphi_1 \vee \mathbf{X}(\varphi_1 \text{ R } \varphi_2))$$

However, if φ_1 is an alternating formula we apply the relation $\varphi_1 \equiv \mathbf{X}\varphi_1$ to obtain the following identities:

$$\varphi_1 \text{ U } \varphi_2 \equiv \varphi_2 \vee (\mathbf{X}\varphi_1 \wedge \mathbf{X}(\varphi_1 \text{ U } \varphi_2)) \quad \varphi_1 \text{ R } \varphi_2 \equiv \varphi_2 \wedge (\mathbf{X}\varphi_1 \vee \mathbf{X}(\varphi_1 \text{ R } \varphi_2))$$

Using these identities, the formula φ_1 is effectively suspended and checked one step later. Similarly, in the case of disjunction or conjunction, each disjunct or conjunct corresponding to an alternating formula is suspended for one step as well. Correctness of these changes clearly follows from properties of alternating

formulae. Note that δ is defined over formulae in positive normal form only. The translation treats each formula $F\psi$ as $ttU\psi$ and each formula $G\psi$ as $(\neg tt)R\psi$.

We introduce further changes to the transition function δ in order to generate automata which exhibits more determinism. In particular, we build a VWAA with only one initial state. Similarly, each state corresponding to a formula of a type $X\varphi$ generates only one successor corresponding to φ . These changes can add an extra initial state and an extra state for each X -subformula comparing to the original construction. However, this drawback is often suppressed due to the consecutive optimizations during the construction of a TGBA.

Now we present a modified construction of VWAA. Given an input LTL formula φ in positive normal form, an equivalent VWAA with a co-Büchi acceptance condition is constructed as $\mathcal{A}_\varphi = (Q, \Sigma, \delta, I, F)$, where Q is the set of all subformulae of φ , Σ and F are defined as in the original construction, $I = \{\varphi\}$, and δ is defined as follows:

$$\begin{aligned} \delta(tt) &= \{(\Sigma, \emptyset)\} \\ \delta(p) &= \{(\Sigma_p, \emptyset)\} \text{ where } \Sigma_p = \{a \in \Sigma \mid p \in a\} \\ \delta(\neg p) &= \{(\Sigma_{\neg p}, \emptyset)\} \text{ where } \Sigma_{\neg p} = \Sigma \setminus \Sigma_p \\ \delta(X\psi) &= \{(\Sigma, \{\psi\})\} \\ \delta(\psi_1 \vee \psi_2) &= \Delta(\psi_1) \cup \Delta(\psi_2) \\ \delta(\psi_1 \wedge \psi_2) &= \Delta(\psi_1) \otimes \Delta(\psi_2) \\ \delta(\psi_1 U \psi_2) &= \begin{cases} \Delta(\psi_2) \cup (\{(\Sigma, \{\psi_1\})\} \otimes \{(\Sigma, \{\psi_1 U \psi_2\})\}) & \text{if } \psi_1 \text{ is alternating,} \\ \Delta(\psi_2) \cup (\Delta(\psi_1) \otimes \{(\Sigma, \{\psi_1 U \psi_2\})\}) & \text{otherwise.} \end{cases} \\ \delta(\psi_1 R \psi_2) &= \begin{cases} \Delta(\psi_2) \otimes (\{(\Sigma, \{\psi_1\}), (\Sigma, \{\psi_1 R \psi_2\})\}) & \text{if } \psi_1 \text{ is alternating,} \\ \Delta(\psi_2) \otimes (\Delta(\psi_1) \cup \{(\Sigma, \{\psi_1 R \psi_2\})\}) & \text{otherwise.} \end{cases} \\ \Delta(\psi) &= \begin{cases} \{(\Sigma, \{\psi\})\} & \text{if } \psi \text{ is a temporal alternating formula,} \\ \delta(\psi) & \text{if } \psi \text{ is a temporal formula that is not alternating.} \end{cases} \\ \Delta(\psi_1 \vee \psi_2) &= \Delta(\psi_1) \cup \Delta(\psi_2) \\ \Delta(\psi_1 \wedge \psi_2) &= \Delta(\psi_1) \otimes \Delta(\psi_2) \end{aligned}$$

Motivation for our changes in the translation can be found in Figures 1 and 2. Each figure contains (a) the VWAA constructed by the original translation and (b) the VWAA constructed by our translation with suspension. Figure 1 shows the effect of suspension of alternating subformula Gfa in computation of transitions leading from the initial state. It can be easily proved that whenever one start with a formula reduced according to Section 4, then each suspension of an alternating temporal subformula leads just to reduction of transitions in the resulting VWAA, i.e., no state is added. On the other hand, if an alternating non-temporal subformula ψ is suspended or the new definition of $\delta(X\psi)$ is used, then the resulting VWAA can contain one more reachable state corresponding to the formula ψ . However, other states may become unreachable and, in particular, the automaton can also have more deterministic states as illustrated by Figure 2.

Optimization of VWAA. In the original algorithm, the VWAA is optimized before it is translated to a TGBA. In particular, if there are two transitions




Fig. 1. VWAA for $(GFa) U b$ generated by (a) the translation of [11], (b) our translation with suspension, and (c) our translation with suspension and further determinization




Fig. 2. VWAA for $X(a \vee b)$ generated by (a) the translation of [11], (b) our translation with suspension, and (c) the translation with suspension and further determinization

$t_1 = (q, \alpha_1, O_1)$ and $t_2 = (q, \alpha_2, O_2)$ satisfying $\alpha_2 \subseteq \alpha_1$ and $O_1 \subseteq O_2$, then t_2 is removed as it is implied by t_1 .

We suggest a generalization of this principle: if $O_1 \subsetneq O_2$ then replace the label α_2 in t_2 by $\alpha_2 \wedge \neg \alpha_1$. If $O_1 = O_2$, replace both transitions by the transition $(q, \alpha_1 \vee \alpha_2, O_1)$. Note that if $\alpha_2 \Rightarrow \alpha_1$, i.e. $\alpha_2 \subseteq \alpha_1$, then $\alpha_2 \wedge \neg \alpha_1 \equiv \neg tt$ and transition t_2 can be removed as before. Our generalized optimization rule increase determinism of the produced VWAA as illustrated by automata (c) of Figures 1 and 2.

6 Improvements in VWAA to TGBA Translation

First, we recall the translation of VWAA to TGBA introduced in [11]. Let $\mathcal{A}_\varphi = (Q, \Sigma, \delta, I, F)$ be a VWAA with a co-Büchi acceptance condition. We define $\mathcal{G}_\mathcal{A} = (Q', \Sigma, \delta', I, T)$ to be a TGBA where:

- $Q' = 2^Q$, i.e. a state is a set of states of \mathcal{A}_φ and represents their conjunction,
- $\delta''(\{q_1, q_2, \dots, q_n\}) = \bigotimes_{i=1}^n \delta(q_i)$ is the non-optimized transition function,
- δ' is the optimized transition function defined as the set of \preceq -minimal transitions of δ'' where the relation \preceq is defined by $t_1 \preceq t_2$ iff $t_1 = (O, \alpha_1, O_1)$, $t_2 = (O, \alpha_2, O_2)$, $\alpha_2 \subseteq \alpha_1$, $O_1 \subseteq O_2$, and $\forall T_f \in \mathcal{T}, t_2 \in T_f \Rightarrow t_1 \in T_f$, and
- $\mathcal{T} = \{T_f \mid f \in F\}$ where $T_f = \{(O, \alpha, O') \mid f \notin O' \text{ or } \exists(\beta, O'') \in \delta(f), \alpha \subseteq \beta \text{ and } f \notin O'' \subseteq O'\}$.

Improved Translation. Our algorithm for a VWAA to TGBA translation differs from the original one only in definition of δ , where we also integrate the idea of suspension of alternating formulae. Recall that each state q_i of a VWAA is a subformula of an input LTL formula and each state of a TGBA is identified with a conjunction of states of a VWAA. Let $O = \{q_1, \dots, q_n\}$ be a state of a TGBA. Then transitions leading from O in a TGBA correspond to combinations of transitions leading from q_1, \dots, q_n in a VWAA. If q_i is an alternating formula and thus it satisfies $q_i \equiv Xq_i$, we can effectively decrease the number of transition combinations that need to be considered during computation of $\delta'(O)$ provided we suspend a full processing of q_i to the succeeding states of the TGBA. More precisely, for the purpose of computation of $\delta'(O)$, we set $\delta(q_i) = \{(\Sigma, \{q_i\})\}$. To construct a TGBA equivalent to the VWAA, we have to ensure that q_i will not be suspended forever during any accepting run of the TGBA. Hence, we enable suspension only in the states that are not on any accepting cycle in a TGBA.

Let M be the minimal set containing all VWAA states of the form $\psi R \rho$ and all subformulae of their right operands ρ . One can easily observe each TGBA state lying on some accepting cycle is a subset of M . The VWAA states outside M , called *progress formulae*, push TGBA computations towards accepting cycles. Suspension is enabled in a TGBA state only if it contains a progress formula. However, if all progress formulae in a TGBA state are alternating, their suspension is not allowed (as suspended progress formulae would not enforce any progress).

Formally, for each TGBA state $O = \{q_1, \dots, q_n\}$ we define $\delta''(O)$ as follows:

$$\delta''(O) = \bigotimes_{i=1}^n \delta_O(q_i), \text{ where}$$

$$\delta_O(q_i) = \begin{cases} \{(\Sigma, \{q_i\})\} & \text{if } O \text{ contains a progress non-alternating formula} \\ & \text{and } q_i \text{ is an alternating formula,} \\ & \text{or } O \text{ contains a progress formula} \\ & \text{and } q_i \text{ is an alternating non-progress formula,} \\ \delta(q_i) & \text{otherwise.} \end{cases}$$

We have obtained better results when we restrict the definition of progress formulae to temporal progress formulae.

Note that the original translation of VWAA to TGBA uses a correct but nonstandard definition of accepting sets T_f . In fact, our modification is correct




Fig. 3. A VWAA \mathcal{A}_ψ corresponding to $GFa \wedge Fb$




Fig. 4. A TGBA \mathcal{G}_ψ corresponding to the VWAA of Figure 3




Fig. 5. Transitions leading from state $\{1, 2\}$ in the TGBA constructed from the VWAA of Figure 3 by (a) the translation of [11] and by (b) our translation with suspension

only if we change the definition of these sets to the natural one (see [1] for an explanation). Intuitively, for each accepting state f of the VWAA with a co-Büchi acceptance, we compute a set T_f of all TGBA transitions that do not contain any VWAA transition looping in f . Formally, $\mathcal{T} = \{T_f \mid f \in F\}$ where

$$T_f = \{(O, \alpha, O') \mid f \notin O \text{ or } (\exists(\beta, O'') \in \delta(f), \exists(\gamma, O''') \in \bigotimes_{f' \in O \setminus \{f\}} \delta(f')) \text{ such that } f \notin O'', \alpha = \beta \wedge \gamma, \text{ and } O' = O'' \cup O'''\}.$$

To demonstrate the effect of suspension during the construction of a TGBA, consider the VWAA \mathcal{A}_ψ for the formula $\psi = GFa \wedge Fb$ depicted in Figure 3. The construction of an equivalent TGBA \mathcal{G}_ψ starts in the initial state $\{1, 2\}$ that corresponds to a conjunction of states 1 and 2 of \mathcal{A}_ψ . Figure 5 depicts the transitions of \mathcal{G}_ψ leading from the initial state when constructed by (a) the original translation of [11] and by (b) our translation with suspension. Note that the state 1 corresponding to the alternating formula GFa is suspended in the TGBA state $\{1, 2\}$ as the state 2 corresponds to a non-alternating progress formula Fb . In both cases, the TGBA has two sets of accepting transitions, T_2 and T_3 . Each transition in the TGBA is labelled by a propositional formula over AP and by a subset of $\{2, 3\}$ indicating to which sets of T_2, T_3 the transition belongs.

Comparing to the original VWAA to TGBA translation without any optimizations, the application of suspension leads to automata with fewer states.

However, if we enable the optimizations suggested in [11], the original translation often constructs automata with the same number of states as our translation with suspension. For example, in the TGBA constructed from the VWAA of Figure 3, the optimizations merge states $\{1, 2, 3\}$ and $\{1, 3\}$ with $\{1, 2\}$ and $\{1\}$, respectively. In this particular case, both approaches lead to the same automaton \mathcal{G}_ψ as shown in Figure 4. However, this is not the case in general. Using suspension, automata with either more or less states can be constructed. However, the translation with suspension is usually slightly faster.

In addition, we detect that both the original and the improved algorithms spend a lot of time when computing transitions of TGBA states equivalent to a formula of the form $\rho = G\alpha_0 \wedge \bigwedge_{1 \leq i \leq n} GF\alpha_i$ where $n \geq 0$ and $\alpha_0, \alpha_1, \dots, \alpha_n$ are formulae without any temporal operator. As such TGBA states are very frequent in practice, we use an optimization that detects these TGBA states and directly constructs the optimal transitions.

7 Optimization of BA

We slightly modify one optimization rule suggested in [11]. It is applied on a resulting BA. The rule says that states q_1 and q_2 of a BA can be merged if $\delta(q_1) = \delta(q_2)$ and $q_1 \in F \iff q_2 \in F$. This rule typically fails to merge the states with a self loop. We suggest to add a new rule where the condition $\delta(q_1) = \delta(q_2)$ is replaced by $\delta(q_1)[q_1/r] = \delta(q_2)[q_2/r]$, where r is a fresh artificial state and $\delta(q)[q/r]$ is a $\delta(q)$ with all occurrences of q as a target node replaced by r .

8 Implementation and Experimental Result

We have implemented all the modifications suggested in the previous sections (and formula reduction rules suggested in [9]) in order to evaluate their effect. The implementation is based on LTL2BA and therefore called *LTL3BA*. Besides the changed algorithms, we also made some other, implementation related changes. In particular, we represent transition labels by BDDs and transitions are represented by C++ STL containers.

In this section, we compare LTL3BA with LTL2BA (v1.1) and SPOT (v0.7.1). For the comparison of results, we use `lbt` testbench tool [19] to measure, for each translator, the number of states and transitions¹ of resulting automata, and the time of the computation. Further, we extend `lbt` to count the number of produced deterministic automata. To be able to compare the results, we set SPOT (option `-N`) to output automata in the form of never claim for SPIN as that is the output of LTL2BA as well. All experiments were done on a server with 8 processors Intel[®] Xeon[®] X7560, 448 GiB RAM and a 64-bit version of GNU/Linux.

¹ To solve the problem with different representation of transitions in automata produced by different tools, we count all transitions leading from a state q to a state r as one.

Table 1. Comparison of translators on two sets of random formulae. Time is in seconds, 'det. BA' is the number of deterministic automata produced by the translator. Note that, using WDBA minimization, SPOT failed to translate 6 formulae of Benchmark2 within the one hour limit. In order to see the effect of WDBA minimization to other formulae, we state in braces the original results increased by the values obtained when these 6 formulae were translated without WDBA minimization.

Translator	Benchmark1				Benchmark2			
	States	Trans.	Time	det. BA	States	Trans.	Time	det. BA
SPOT	1 561	5 729	7.47	55	14 697	95 645	68.46	221
SPOT+WDBA	1 587	5 880	10.81	88	13 097 (14 408)	77 346 (94 248)	5 916.45 (5 919.43)	373 (373)
LTL2BA	2 118	9 000	0.81	25	24 648	232 400	18.57	84
LTL3BA(1)	1 621	5 865	1.26	27	17 107	129 774	22.25	92
LTL3BA(1,2)	1 631	6 094	1.41	54	15 936	115 624	9.04	237
LTL3BA(1,2,3)	1 565	5 615	1.41	54	14 113	91 159	8.53	240
LTL3BA(1,2,3,4)	1 507	5 348	1.38	54	13 244	85 511	8.30	240

First we compare the translators on two sets, Benchmark1 and Benchmark2, of random formulae generated by `lbt`. Benchmark1 contains 100 formulae of the length 15–20 and their negations. Benchmark2 contains 500 formulae of the length 15–30 and their negations. The exact `lbt` parameters used to generate the formulae are in [1]. Table 1 presents the cumulative results of translations of all formulae in the two sets. The table also illustrates the gradual effect of modifications of each step of the translation (1,2,3,4 refers to modifications introduced in Sections 4, 5, 6, and 7 in the respective order; e.g. LTL3BA(1) uses the original algorithm with our formula reduction while LTL3BA(1,2,3,4) refers to the translation with all the suggested modifications). Finally, the table contains the results for SPOT with WDBA minimization, which has the longest running time but provides the best results. The automata produced by LTL3BA are in sum slightly better than the automata produced by SPOT. Further, LTL3BA seems to be much faster.

Further, we compare the execution time of translators running on parametric formulae from [11] and [16]. We use SPOT with the recommended option `-r4`, i.e. with the input formula reduction as the only optimization. To get a comparable settings of LTL3BA, we switched off the generalized optimization of VWAA. We gradually increase the parameter of the formulae until a translator fails to finish the translation in one hour limit. The results are partly depicted in Figure 6 (the rest is in [1]).

The graphs show that, in general, LTL3BA is slightly slower than LTL2BA and faster than SPOT on small formulae. With increasing parameter, LTL3BA outperforms LTL2BA (with exception of $S(n)$ where LTL2BA fails before its running time reaches the limit), while SPOT sometimes remains slower, but sometimes eventually outperform LTL3BA.

For more experimental results (including the benchmark of [2]) see [1].

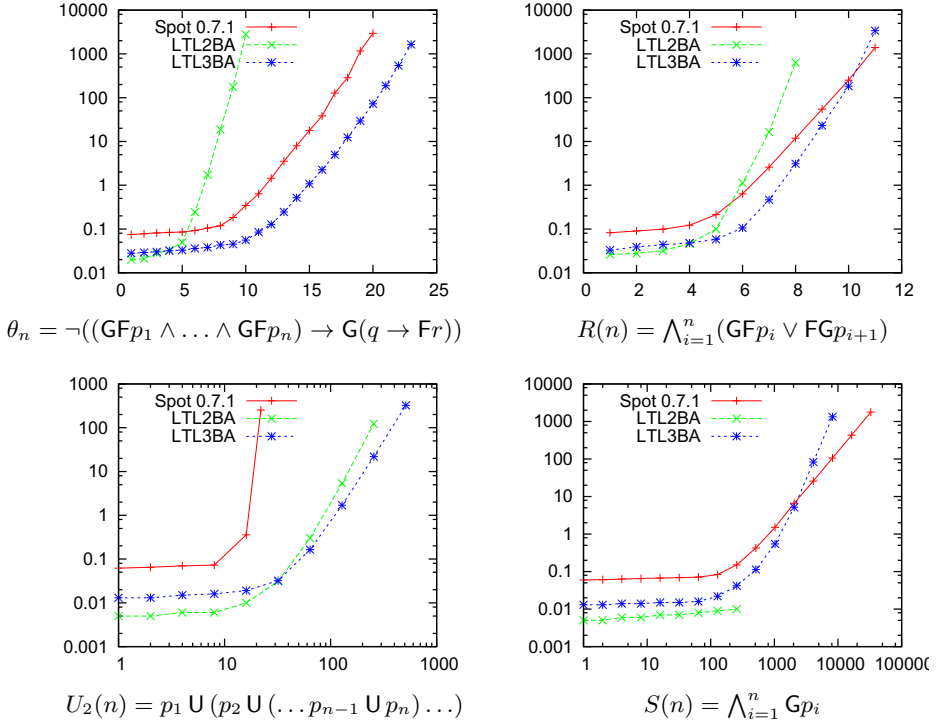


Fig. 6. Time consumption for parametric formulae constructed within an hour (the vertical axes are logarithmic and represent time in seconds, while the horizontal axes are linear or logarithmic and represent the parameter n)

9 Conclusion

We have focused on LTL to BA translations with the stress on their speed-up while maintaining outputs of a good quality. We have introduced several modifications of LTL2BA on both algorithmic and implementation levels. Among others, we have identified an LTL subclass of “alternating” formulae, validity of which does not depends on any finite prefix of the word.

Our experimental results indicate that our modifications have a mostly positive effect on each step of the translation. The new translator called LTL3BA is usually faster than the original LTL2BA and it produces smaller and more deterministic automata. Moreover, comparison of LTL3BA and the current version of SPOT (run without WDBA minimization that is very slow) shows that the produced automata are of similar quality and LTL3BA is usually faster.

LTL3BA has served as an experimental tool to demonstrate that our modifications are improvements and their applicability to other LTL to BA translations is a subject of further research.

Acknowledgments. The authors would like to thank three anonymous referees and Alexandre Duret-Lutz for valuable comments.

References

1. Babiak, T., Křetínský, M., Řehák, V., Strejček, J.: LTL to Büchi Automata Translation: Fast and More Deterministic. CoRR, abs/1201.0682 (2012)
2. Cichoń, J., Czubak, A., Jasiński, A.: Minimal Büchi automata for certain classes of LTL formulas. In: DEPCOS-RELCOMEX 2009, pp. 17–24. IEEE (2009)
3. Couvreur, J.-M.: On-the-Fly Verification of Linear Temporal Logic. In: Wing, J.M., Woodcock, J. (eds.) FM 1999. LNCS, vol. 1708, pp. 253–271. Springer, Heidelberg (1999)
4. Daniele, M., Giunchiglia, F., Vardi, M.Y.: Improved Automata Generation for Linear Temporal Logic. In: Halbwachs, N., Peled, D.A. (eds.) CAV 1999. LNCS, vol. 1633, pp. 249–260. Springer, Heidelberg (1999)
5. Dax, C., Eisinger, J., Klaedtke, F.: Mechanizing the Powerset Construction for Restricted Classes of ω -Automata. In: Namjoshi, K.S., Yoneda, T., Higashino, T., Okamura, Y. (eds.) ATVA 2007. LNCS, vol. 4762, pp. 223–236. Springer, Heidelberg (2007)
6. Duret-Lutz, A.: LTL translation improvements in Spot. In: VECoS 2011, eWiC. British Computer Society (2011) (to appear)
7. Duret-Lutz, A., Poitrenaud, D.: SPOT: An extensible model checking library using transition-based generalized Büchi automata. In: MASCOTS 2004, pp. 76–83. IEEE (2004)
8. Ehlers, R., Finkbeiner, B.: On the Virtue of Patience: Minimizing Büchi Automata. In: van de Pol, J., Weber, M. (eds.) Model Checking Software. LNCS, vol. 6349, pp. 129–145. Springer, Heidelberg (2010)
9. Etesami, K., Holzmann, G.J.: Optimizing Büchi Automata. In: Palamidessi, C. (ed.) CONCUR 2000. LNCS, vol. 1877, pp. 153–167. Springer, Heidelberg (2000)
10. Fritz, C.: Constructing Büchi Automata from Linear Temporal Logic Using Simulation Relations for Alternating Büchi Automata. In: Ibarra, O.H., Dang, Z. (eds.) CIAA 2003. LNCS, vol. 2759, pp. 35–48. Springer, Heidelberg (2003)
11. Gastin, P., Oddoux, D.: Fast LTL to Büchi Automata Translation. In: Berry, G., Comon, H., Finkel, A. (eds.) CAV 2001. LNCS, vol. 2102, pp. 53–65. Springer, Heidelberg (2001)
12. Giannakopoulou, D., Lerda, F.: From States to Transitions: Improving Translation of LTL Formulae to Büchi Automata. In: Peled, D.A., Vardi, M.Y. (eds.) FORTE 2002. LNCS, vol. 2529, pp. 308–326. Springer, Heidelberg (2002)
13. Löding, C.: Efficient minimization of deterministic weak omega-automata. Information Processing Letters 79(3), 105–109 (2001)
14. Manna, Z., Pnueli, A.: A hierarchy of temporal properties. In: PODC 1990, pp. 377–410. ACM Press (1990)
15. Pnueli, A.: The temporal logic of programs. In: FOCS 1977, pp. 46–57. IEEE (1977)
16. Rozier, K.Y., Vardi, M.Y.: LTL Satisfiability Checking. In: Bošnački, D., Edelkamp, S. (eds.) SPIN 2007. LNCS, vol. 4595, pp. 149–167. Springer, Heidelberg (2007)
17. Sebastiani, R., Tonetta, S.: “More Deterministic” vs. “Smaller” Büchi Automata for Efficient LTL Model Checking. In: Geist, D., Tronci, E. (eds.) CHARME 2003. LNCS, vol. 2860, pp. 126–140. Springer, Heidelberg (2003)
18. Somenzi, F., Bloem, R.: Efficient Büchi Automata from LTL Formulae. In: Emerson, E.A., Sistla, A.P. (eds.) CAV 2000. LNCS, vol. 1855, pp. 248–263. Springer, Heidelberg (2000)
19. Tauriainen, H., Heljanko, K.: Testing LTL formula translation into Büchi automata. STTT 4(1), 57–70 (2002)