# LTL translation improvements in Spot 1.0

## Alexandre Duret-Lutz

EPITA's Research and Development Laboratory (LRDE),
14-16 rue Voltaire,
94270 Le Kremlin-Bicêtre, France
E-mail: adl@lrde.epita.fr

**Abstract:** Spot is a library of model-checking algorithms started in 2003. This paper focuses on its module for translating linear-time temporal logic (LTL) formulas into Büchi automata: one of the steps required in the automata-theoretic approach to LTL model-checking. We detail the different algorithms involved in this translation: the core translation itself, which performs many simplifications thanks to its use of binary decision diagrams; the pre-processing of the LTL formulas with rewriting rules chosen to help their translation; and various post-processing algorithms whose use depends on the intent of the translation: do we favour deterministic automata, or small automata? Using different benchmarks, we show how Spot competes with other LTL translators, and how it has improved over the years.

**Keywords:** formal methods; model checking; Büchi automata; linear-time temporal logic; LTL; temporal logic; translation; simplifications; software; verification; implementation.

**Biographical notes:** Alexandre Duret-Lutz is a Researcher and Teacher at EPITA, a private computer science school near Paris. He maintains Spot, a model checking library that he started during is PhD, and is passionate about LTL translation.

This paper is a revised and expanded version of a paper entitled 'LTL translation improvement in Spot' presented at the 5th Int. Workshop on Verification and Evaluation of Computer and Communication Systems (VECoS'11), Tunis, September 2011.

## 1 Introduction

One of the first steps of the automata-theoretic approach to model checking of linear-time properties (Vardi, 1996, 2007) is to translate the property to verify into an $\omega$-automaton. This automaton is then synchronised with a model of the system in order to find executions that invalidate the property. By constructing a smaller or more

deterministic property automaton, we can hope (this is generally the case) to obtain a smaller synchronised product to explore, resulting in faster model checking.

The Spot library (Duret-Lutz and Poitrenaud, 2004) offers algorithms to realise the above automata-theoretic approach. A salient feature of Spot is its preference for using transition-based generalised Büchi automata (TGBA) instead of the more commonly used Büchi automata (BA). Section 2 explains the difference.

This paper attempts to give a global view of the different algorithms involved into the LTL-to-TGBA or LTL-to-BA translation of Spot, to explain why it often produces smaller automata than other available translators, and why it does not always produce them as fast. Along the way, we point some steps (like the degeneralisation) that could probably be improved. We believe the insight we provide into the implementation of Spot should be helpful to anyone devising a new translator.

Spot actually offers four translation procedures, and we shall only discuss the most efficient one, derived from an algorithm by Couvreur (1999).

A previous version of this paper was presented at VECOS '11 (Duret-Lutz, 2011). The text has been augmented to discuss new optimisations implemented between Spot 0.7 and Spot 1.0, and presents new benchmarks featuring more linear-time temporal logic (LTL) translators.

We assume the reader is familiar with LTL (Clarke et al., 2000) and binary decision diagrams (Bryant, 1986), abbreviated as BDDs in the sequel.

This paper is organised as follows. Section 2 defines TGBA as opposed to BA. Section 3 presents the core of the translation algorithm, with an emphasis on the optimisations that are enabled by the use of BDDs, and discusses some improvements to this translation. In Sections 4 and 5 we discuss pre-processing and post-processings. Finally, Section 6 compares Spot with other translators on various benchmarks.

Throughout the paper, the reader is invited to play with an on-line version of the translator at http://spot.lip6.fr/ltl2tgba.html. This page has options for many optimisations discussed herein.

## 2   Two kinds of Büchi automata

Let *AP* be a set of *atomic propositions*, i.e., propositional variables that may be true or false in the system. $2^{AP}$ denotes the set of minterms (or assignments) over *AP*, and $2^{2^{AP}}$, interpreted as the set of sums of minterms, denotes the Boolean formulas over *AP*.

*Definition 1:* A Büchi automaton is a tuple $\mathcal{B} = \langle AP, \mathcal{Q}, q^0, \mathcal{F}, \delta \rangle$ where *AP* is a set of *atomic propositions*, $\mathcal{Q}$ is a finite set of states, $q^0 \in \mathcal{Q}$ is the initial state, $\mathcal{F} \subseteq \mathcal{Q}$ is a set of acceptance states, and $\delta \subseteq \mathcal{Q} \times 2^{AP} \times \mathcal{Q}$ is a transition relation in which each transition is labelled by a Boolean assignment.
An infinite word $c_0 c_1 c_2 \ldots \in (2^{AP})^\omega$ of assignments is *accepted* by $\mathcal{B}$ if there exists a run of $\mathcal{A}$, say $(q^0, l_0, q_1)(q_1, l_1, q_2)(q_2, l_2, q_3)\ldots \in \delta^\omega$, that recognises the word ($\forall i, c_i = l_i$) and that visits infinitely many acceptance states ($\forall i \geq 0, \exists j \geq i, q_j \in \mathcal{F}$).

A common implementation technique is to group transitions with common source and destination into edges labelled by Boolean formulas. E.g., the three *transitions* $(q_1, a\bar{b}, q_2)$, $(q_1, ab, q_2)$, and $(q_1, \bar{a}b, q_2)$ can be represented by one *edge* $(q_1, a \vee b, q_2)$.

A transition-based generalised Büchi automaton (TGBA) is a Büchi automaton in which multiple acceptance marks are carried by the transitions.

*Definition 2:* A TGBA is a tuple $\mathcal{T} = \langle AP, \mathcal{Q}, q^0, \mathcal{F}, \delta \rangle$ where $AP$ is a set of *atomic propositions*, $\mathcal{Q}$ is a finite set of states, $q^0 \in \mathcal{Q}$ is the initial state, $\mathcal{F} = \{f_1, f_2, \ldots, f_n\}$ is a finite set of *acceptance marks*, $\delta \subseteq \mathcal{Q} \times 2^{AP} \times 2^{\mathcal{F}} \times \mathcal{Q}$ is a transition relation in which each transition is labelled by a Boolean assignment and a set of acceptance marks.

An infinite word $c_0 c_1 c_2 \ldots \in (2^{AP})^\omega$ of assignments is *accepted* by $\mathcal{T}$ if there exists a run of $\mathcal{A}$, say $(q^0, l_0, F_0, q_1)(q_1, l_1, F_1, q_2)(q_2, l_2, F_2, q_3) \ldots \in \delta^\omega$, that recognises the word ($\forall i, c_i = l_i$) and that visits each acceptance mark infinitely often ($\forall f \in \mathcal{F}$, $\forall i \geq 0, \exists j \geq i, f \in F_j$).

Similarly, transitions that share the same source, destination and acceptance mark may be implemented by a single edge labelled by a Boolean formula. For simplicity, we only display these edges on the figures.

Figure 1 illustrates these definitions with two automata that recognise the LTL property: $\mathsf{G}\,\mathsf{F}\,a \wedge \mathsf{G}\,\mathsf{F}\,b$. The infinite sequence $\begin{smallmatrix} a: & 1 & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 \\ b: & 0 & 0 & 1 & 0 & 0 & 1 & 0 & 0 & 1 \end{smallmatrix} \ldots$ will be accepted by $\mathcal{T}_1$ because it visits the top and right loops infinitely often, therefore all acceptance marks are seen infinitely often. Similarly this sequence visits the only acceptance state of $\mathcal{B}_1$ infinitely often.

**Figure 1** Two automata recognising the LTL formula $\mathsf{G}\,\mathsf{F}\,a \wedge \mathsf{G}\,\mathsf{F}\,b$
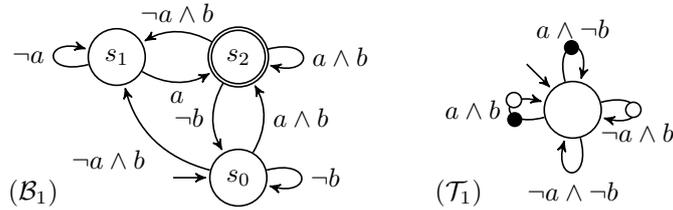


Notes: $\mathcal{B}_1$: Büchi automaton with a single acceptance state (double circle). $\mathcal{T}_1$: TGBA with $\mathcal{F} = \{\circ, \bullet\}$.

Spot is built around TGBAs and can perform the entire model-checking approach with these automata. However most other model-checking tools use BA. Fortunately, TGBAs can be *degeneralised* into BA by an operation discussed in Section 5.4. Automaton $\mathcal{B}_1$ in Figure 1 was obtained by degeneralising $\mathcal{T}_1$.

We will often name the states of automata with the LTL formula they accept. These extra annotations have no influence on the behaviour of the automata.

In a Büchi automaton, we say that a strongly connected component (SCC) is accepting if it contains some accepting state. In a TGBA an SCC is accepting if for each acceptance mark it contains at least one marked transition.

## 3    From LTL to TGBA

The algorithm of Couvreur (1999) for the translation of LTL automata into TGBA is based on a tableau method. Although the following explanations are self-contained, we refer the reader to Duret-Lutz and Poitrenaud (2004) for an illustration of this algorithm as a tableau that can be used to build generalised BA with state-based or transition-based acceptance conditions. Here we shall present the algorithm at a lower level to explain how the use of BDDs helps the translation.

To put this algorithm in context, the complete translation procedure to go from LTL to a Büchi automaton can be presented as four steps:

1    Simplify the LTL formula syntactically, e.g., rewrite $\mathsf{F}\,\mathsf{F}\,a$ (a 3-state automaton) into $\mathsf{F}\,a$ (2 states). These pre-processings are discussed in Section 4.

2    Translate the simplified formula into a TGBA using the algorithm presented in this section.

3    Post-process the resulting TGBA, e.g., by pruning useless SCCs, or running various simulation-based reductions or minimisations discussed in Section 5.

4    If desired (and needed after the previous post-processing) degeneralise the TGBA into a Büchi automaton, as discussed in Section 5.4.

### 3.1    Basic translation

If we omit BDDs, the procedure is simple enough to be performed by hand on a paper or blackboard. The algorithm generates an automaton whose states corresponds to LTL formulas. The initial state is the formula to translate. This formula is then rewritten as a sum of products where the only temporal operator allowed at the top level is $\mathsf{X}$.

For instance if we were to translate $\Psi = (\mathsf{X}\,a) \wedge (b\,\mathsf{U}\,\neg a)$ we would use the fact that $\varphi\,\mathsf{U}\,\psi = \psi \vee (\varphi \wedge \mathsf{X}(\varphi\,\mathsf{U}\,\psi))$ to rewrite $\Psi$ as $(\neg a \wedge \mathsf{X}\,a) \vee (b \wedge \mathsf{X}\,a \wedge \mathsf{X}(b\,\mathsf{U}\,\neg a))$. Reading this formula, it is clear that a state that must recognise $\Psi$ should either accept an assignment compatible with $\neg a$ and verify $a$ at the next step, or accept an assignment compatible with $b$ and then verify $a \wedge (b\,\mathsf{U}\,\neg a)$ at the next step. The start of the automaton is thus as follows:
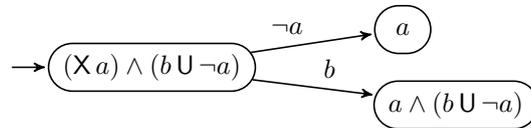


The procedure should then be applied similarly on the new states. There is little technicality that has to be taken into account when translating the $\varphi\,\mathsf{U}\,\psi$ operator: the formula $\psi$ *must* be satisfied eventually, it cannot be postponed continuously. This is solved in the translation by making a *promise to fulfill* $\psi$ while rewriting the formula. The actual rewriting rule used for $\mathsf{U}$ is: $\varphi\,\mathsf{U}\,\psi = \psi \vee (\varphi \wedge \mathsf{X}(\varphi\,\mathsf{U}\,\psi) \wedge \mathsf{P}\,\psi)$, with the operator $\mathsf{P}$ denoting an explicit promise.

All these formulas can be simplified using classical Boolean rules like $(\alpha \wedge \beta) \vee \alpha = \alpha$ to kill some terms (even $\mathsf{X}\,\varphi$ or $\mathsf{P}\,\varphi$). This is where using BDD really helps. The core of the translation is the rewriting function $r(f)$ defined recursively as in Figure 2. It encodes outgoing transitions using BDD variables of the form $\mathrm{Var}[p]$, $\mathrm{Nxt}[f]$, $\mathrm{P}[f]$, created as needed to represent respectively atomic propositions, $\mathsf{X}\,f$ formulas, and $\mathsf{P}\,f$ promises. The given definition assumes that the LTL formula is specified into *negative normal form*, where negation operators appear only in front of atomic propositions.

Applying $r$ on our example, we obtain:

$$
\begin{aligned}
r((\mathsf{X}\,a) \wedge (b\,\mathsf{U}\,\neg a)) &= r(\mathsf{X}\,a)\ \wedge\ r(b\,\mathsf{U}\,\neg a) \\
&= \mathrm{Nxt}[a]\ \wedge\ (r(\neg a)\ \vee\ (r(b)\ \wedge\ \mathrm{Nxt}[b\,\mathsf{U}\,\neg a] \wedge \mathrm{P}[\neg a])) \\
&= \mathrm{Nxt}[a]\ \wedge\ (\neg\mathrm{Var}[a]\ \vee\ (\mathrm{Var}[b] \wedge \mathrm{Nxt}[b\,\mathsf{U}\,\neg a] \wedge \mathrm{P}[\neg a])) \\
&= (\neg\mathrm{Var}[a] \wedge \mathrm{Nxt}[a])\ \vee\ (\mathrm{Var}[b] \wedge \mathrm{Nxt}[a] \wedge \mathrm{Nxt}[b\,\mathsf{U}\,\neg a] \wedge \mathrm{P}[\neg a])
\end{aligned}
$$

Which corresponds to:
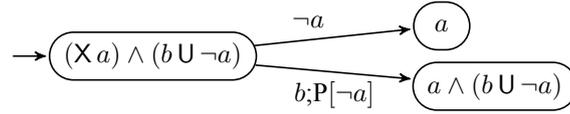


**Figure 2**  Recursive rules to translate an LTL formula into a BDD

$$
\begin{aligned}
&& r(\mathsf{X}\,f) &= \mathrm{Nxt}[f] \\
r(\top) &= \top & r(\mathsf{F}\,f) &= r(f)\ \vee\ (\mathrm{Nxt}[\mathsf{F}\,f] \wedge \mathrm{P}[f]) \\
r(\bot) &= \bot & r(\mathsf{G}\,f) &= r(f)\ \wedge\ \mathrm{Nxt}[\mathsf{G}\,f] \\
r(p) &= \mathrm{Var}[p] & r(f\,\mathsf{U}\,g) &= r(g)\ \vee\ (r(f) \wedge \mathrm{Nxt}[f\,\mathsf{U}\,g] \wedge \mathrm{P}[g]) \\
r(\neg p) &= \neg\mathrm{Var}[p] & r(f\,\mathsf{W}\,g) &= r(g)\ \vee\ (r(f) \wedge \mathrm{Nxt}[f\,\mathsf{W}\,g]) \\
r(f \vee g) &= r(f) \vee r(g) & r(f\,\mathsf{R}\,g) &= r(g)\ \wedge\ (r(f) \vee \mathrm{Nxt}[f\,\mathsf{R}\,g]) \\
r(f \wedge g) &= r(f) \wedge r(g) & r(f\,\mathsf{M}\,g) &= r(g)\ \wedge\ (r(f) \vee (\mathrm{Nxt}[f\,\mathsf{M}\,g] \wedge \mathrm{P}[f]))
\end{aligned}
$$

There are several ways to turn a BDD into a sum of products, but because each term of the sum corresponds to a transition in the automaton, redundant terms should be avoided. Furthermore, $\mathrm{Nxt}[]$ and $\mathrm{P}[]$ variables should never be negated. We compute prime implicants using an algorithm from Minato (1992) to that effect.

The complete translation is shown on Figure 3. This automaton is still not a TGBA because it uses promises instead of acceptance marks. To guarantee that a promise holds, the accepted runs of the automaton should never make promises continuously: in other words for each promise $\mathsf{P}\,\varphi$, accepted runs should visit infinitely many transitions that do not make such a promise.[1] This can be encoded as a TGBA by labelling all transitions that do not make promise $\mathsf{P}\,\varphi$ by an acceptance mark associated to $\varphi$. There will be as many acceptance marks as promises. Figure 4 shows the final TGBA.

The pseudo-code for the complete translation algorithm is shown on Figure 5.

**Figure 3**    Translation of $(\mathsf{X}\,a) \wedge (b\,\mathsf{U}\,\neg a)$ using promises
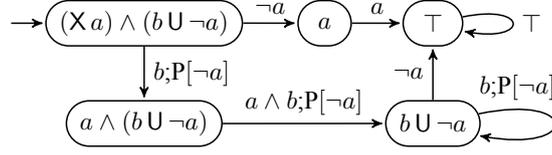


**Figure 4**    Translation of $(\mathsf{X}\,a) \wedge (b\,\mathsf{U}\,\neg a)$ as a TGBA
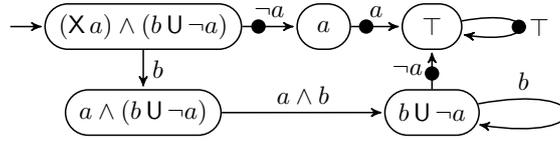


**Figure 5**    Pseudo-code of the algorithm of Couvreur (1999) to translate an LTL formula $f$ into a TGBA

```
ltl_to_tgba_fm(f):
    todo ← {f}; all_acc ← ∅
    a ←new automaton; a.set_initial_state(f)
    while (todo ≠ ∅)
        here ← todo.remove_one()
        forall i in prime_implicants_of(r(here))
            Put i as ⋀_{v∈V} Var[v] ∧ ⋀_{v∈V′} ¬Var[v] ∧ ⋀_{a∈A} P[a] ∧ ⋀_{n∈N} Nxt[n]
            dest ← ⋀_{n∈N} n
            if ¬a.has_state(dest)
                todo.insert(dest)
            a.add_edge(src: here, dst: dest, cond: ⋀_{v∈V} v ∧ ⋀_{v∈V′} ¬v, promises: A)
            all_acc ← all_acc ∪ A
    forall t in a.edges()
        t.acceptance_marks ← all_acc \ t.promises
    return a
```

Note: The function $r(here)$ is defined on Figure 2.

At this point it should be clear that the use of BDDs simplifies every Boolean formulas that label edges. For instance we cannot have an edge labelled by $b \wedge a \wedge \neg b$ because such a conjunction would be simplified by the BDD representation.

Similarly the conversion of the BDD into a sum of prime implicants helps to reduce the number of outgoing arcs of each node.

We experimented with different BDD variable orders, and found it was better to introduce variables in the order they are discovered while applying $r$ recursively.

### 3.2 Using $r$ to identify states

A powerful BDD-based optimisation is to use $r$ to identify some equivalent formulas. Because BDDs have a unique representation, two formulas $\varphi$ and $\psi$ are equivalent if their rewritings are the same BDDs $r(\varphi) = r(\psi)$. The converse does not hold because two equivalent subformulas prefixed with $\mathsf{X}$ might be represented by different $\mathrm{Nxt}[]$ variables. Since $r(\varphi)$ encodes the outgoing edges (labels, promises, and destinations) of the state $\psi$, if $r(\varphi) = r(\psi)$ then the states $\varphi$ and $\psi$ have exactly the same successors and can be merged. Such a reduction occurs when translating $\mathsf{G}\,\mathsf{F}\,a$:

$$r(\mathsf{G}\,\mathsf{F}\,a) = ((\mathrm{Nxt}[\mathsf{F}\,a] \wedge \mathrm{P}[a]) \vee \mathrm{Var}[a]) \wedge \mathrm{Nxt}[\mathsf{G}\,\mathsf{F}\,a]$$

$$r((\mathsf{F}\,a) \wedge \mathsf{G}\,\mathsf{F}\,a) = ((\mathrm{Nxt}[\mathsf{F}\,a] \wedge \mathrm{P}[a]) \vee \mathrm{Var}[a]) \wedge \mathrm{Nxt}[\mathsf{G}\,\mathsf{F}\,a]$$

The result of $r(\mathsf{G}\,\mathsf{F}\,a)$ implies that $\mathsf{G}\,\mathsf{F}\,a$ should have two successors, $\mathsf{G}\,\mathsf{F}\,a$ and $(\mathsf{F}\,a) \wedge \mathsf{G}\,\mathsf{F}\,a$, as shown in the first automaton of Figure 6. However $r((\mathsf{F}\,a) \wedge \mathsf{G}\,\mathsf{F}\,a) = r(\mathsf{G}\,\mathsf{F}\,a)$ so these states can be merged.

**Figure 6** Two translations of $\mathsf{G}\,\mathsf{F}\,a$
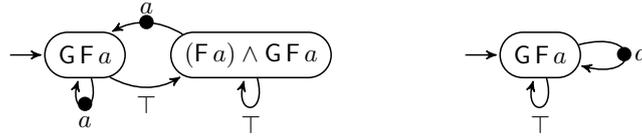


Note: Since $r(\mathsf{G}\,\mathsf{F}\,a) = r((\mathsf{F}\,a) \wedge \mathsf{G}\,\mathsf{F}\,a)$ the two states of the first automaton can be merged, yielding the second automaton.

One way to implement this '$r$-quotienting' automatically is to index the states of the automaton by the BDD $r(\varphi)$ instead of by the LTL formula $\varphi$ (the pseudo-code from Figure 5 does *not* perform this reduction).

This automatic simplification may fail to merge states that have the same successors except for a self-loop because the $\mathrm{Nxt}[]$ variable representing the destination of the self-loop will be different in each state. Babiak et al. (2012) have suggested to improve this case by introducing a unique dummy BDD variable to represent the current state. This optimisation is implemented in their LTL translator, `ltl3ba`, but not yet in Spot.

### 3.3 Better determinism

The determinism of the automata from Figure 6 can be improved using a trick based on the BDD representation of states. Instead of converting the equation $r(\mathsf{G}\,\mathsf{F}\,a) = ((\mathrm{Nxt}[\mathsf{F}\,a] \wedge \mathrm{P}[a]) \vee \mathrm{Var}[a]) \wedge \mathrm{Nxt}[\mathsf{G}\,\mathsf{F}\,a]$ into a sum of products to discover the labels and destinations, we can instead fix one label to discover its destination(s).

Where shall we go if we read $a$? $r(\mathsf{G}\,\mathsf{F}\,a) \wedge \mathrm{Var}[a] = \mathrm{Var}[a] \wedge \mathrm{Nxt}[\mathsf{G}\,\mathsf{F}\,a]$.

If we read $\neg a$? $r(\mathsf{G}\,\mathsf{F}\,a) \wedge \neg\mathrm{Var}[a] = \neg\mathrm{Var}[a] \wedge \mathrm{Nxt}[\mathsf{F}\,a] \wedge \mathrm{P}[a] \wedge \mathrm{Nxt}[\mathsf{G}\,\mathsf{F}\,a]$.

These equations show that all instances of $\top$ in Figure 6 can be replaced by $\neg a$, yielding two deterministic automata.

In an automaton over $n$ atomic propositions ($\mathrm{Var}[a]$, $\mathrm{Var}[b]$,...), there are $2^n$ labels to consider. However the structure of the BDD encoding the formula helps to ignore

useless labels; and in real-world formulas, $n$ is usually small enough to make the enumeration of these labels not perceptible.

While an automaton constructed this way is usually *more* deterministic, it is not necessarily a deterministic automaton. The result of $r(\varphi) \wedge A$ for some $A$ could feature a disjunction, i.e., multiple destinations (the reader is invited to compute $r(\mathsf{F}\,\mathsf{G}\,a) \wedge \mathrm{Var}[a]$ for an example).

In an experiment we translated 92 LTL formulas taken from the literature and compared their translations with and without this optimisation, by synchronising the resulting automata with random state spaces. This technique reduced the number of transitions in the product by 40%, and the number of states by only 0.33%.

### 3.4   Speeding up the translation of G formulas

In his original paper, Couvreur (1999) discussed an optimisation of this translation using a specific rule for formulas of the form $\mathsf{G}\,\mathsf{F}\,f$: $r(\mathsf{G}\,\mathsf{F}\,f) = (r(f) \vee \mathrm{P}[f]) \wedge \mathrm{Nxt}[\mathsf{G}\,\mathsf{F}\,f]$.

This rule avoids the creation of the state $\mathsf{F}\,a \wedge \mathsf{G}\,\mathsf{F}\,a$ during the translation of $\mathsf{G}\,\mathsf{F}\,a$. From a size perspective, it is entirely optional since the $r$-quotienting discussed in Section 3.2 will already identify the two states. However from a time point of view, it is more efficient to construct a single state directly, and avoid many BDD operations. (Spot's translator spends more than half of its run time performing BDD operations.)

We generalised this rule to apply to any subformula that is guaranteed to be repeated in the next state. We modify the G rule of Figure 2 as: $r(\mathsf{G}\,f) = r_\mathsf{G}(f) \wedge \mathrm{Nxt}[\mathsf{G}\,f]$ where $r_\mathsf{G}$ is the recursive function defined by Figure 7. These $r_\mathsf{G}$ rules, called inside $r(\mathsf{G}\,f)$, avoid the creation of the $\mathrm{Nxt}[f]$ variables that would be implied by $\mathrm{Nxt}[\mathsf{G}\,f]$ anyway. In particular, this optimisation halves the time spent translating subformulas of the form $\bigwedge_i \mathsf{G}\,\mathsf{F}\,p_i$ or of the equivalent (but preferred) form $\mathsf{G}\bigwedge_i \mathsf{F}\,p_i$, either of which occur when expressing weak fairness properties.

**Figure 7**   Recursive rules to translate LTL subformulas of G

$$
\begin{aligned}
r_\mathsf{G}(f \wedge g) &= r_\mathsf{G}(f) \wedge r_\mathsf{G}(g) & r_\mathsf{G}(f \,\mathsf{W}\, g) &= r(g) \vee r(f) \\
r_\mathsf{G}(\mathsf{F}\,f) &= r(f) \vee \mathrm{P}[f] & r_\mathsf{G}(f \,\mathsf{R}\, g) &= r_\mathsf{G}(g) \\
r_\mathsf{G}(f \,\mathsf{U}\, g) &= r(g) \vee (r(f) \wedge \mathrm{P}[g]) & r_\mathsf{G}(f \,\mathsf{M}\, g) &= r_\mathsf{G}(g) \wedge (r(f) \vee \mathrm{P}[f]) \\
r_\mathsf{G}(f) &= r(f) \text{ in all other cases}
\end{aligned}
$$

### 3.5   Simplifying promises

Consider the BDD rewriting of $a \,\mathsf{U}\,(b \,\mathsf{U}\, c)$ whose complete automaton is represented with promises on Figure 8:

$r(a \,\mathsf{U}\,(b \,\mathsf{U}\, c))$
$= \mathrm{Var}[c] \vee (\mathrm{Var}[b] \wedge \mathrm{Nxt}[b \,\mathsf{U}\, c] \wedge \mathrm{P}[c]) \vee (\mathrm{Var}[a] \wedge \mathrm{Nxt}[a \,\mathsf{U}\,(b \,\mathsf{U}\, c)] \wedge \mathrm{P}[b \,\mathsf{U}\, c])$

This BDD encodes three transitions, two of which use different promises: $\mathrm{P}[c]$ and $\mathrm{P}[b \,\mathsf{U}\, c]$. However these promises are always issued sequentially: first $\mathrm{P}[b \,\mathsf{U}\, c]$ forbids

runs that continuously stay in the initial state, then if the state $b \cup c$ is reached, $P[c]$ rejects runs that would stay infinitely in that state. In practice, we could have made the same promise, for instance $P[c]$ (the name does not even matter), on all these transitions. If we interpret $P[f]$ as a promise to fulfill $f$ eventually, it is clear that $P[b \cup c]$ and $P[c]$ are two equivalent promises.

**Figure 8** Translation of $a \cup (b \cup c)$ using promises (and without using the determinisation improvement of Section 3.3)
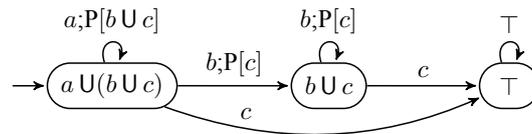


Along these lines, we implement the following simplifications to limit the number of promises introduced: $P[F\,f] = P[f]$, $P[f \cup g] = P[g]$, and $P[f \,M\, g] = P[f]$.

Furthermore, if the top-level formula is a syntactic persistence, only one promise need to be used during the translation and we rewrite any $P[f]$ as $P[\top]$. This optimisation and the class of syntactic persistence formulas are described by Černá and Pelánek (2003).

## 4 Pre-processings

Pre-processing the LTL formula before it is translated helps to speed-up the translation, and to produce smaller automata. Spot distinguishes different kinds of LTL rewritings:

- *Trivial identities* are applied at any time during the construction of a formula (e.g., while they are parsed). These are all based on idempotence of some operators (e.g., $F\,F\,a \equiv F\,a$), or neutral/absorbent operands (e.g., $X \perp \equiv \perp$, $f \wedge \perp \equiv \perp$, $f \wedge \perp \equiv \perp$, etc.).

- *Basic rewritings* are unconditional rewriting rules, such as $G\,X\,f \equiv X\,G\,f$.

- *Eventual and universal rewritings* apply only when some subformulas are *purely universal* Etessami and Holzmann (2000), are *pure eventualities* Etessami and Holzmann (2000), or are what Babiak et al. (2012) have called *alternating* formulas. As an example $F\,G\,F\,a$ can be rewritten as $G\,F\,a$ because the latter is a pure eventuality.

- *Implication-based rewritings* apply only in cases where one subformula can be shown to imply another subformula. For instance under the hypothesis that $f \rightarrow g$, we have $f \cup g \equiv g$. There are two ways to detect such implications: they can be approximated syntactically (Somenzi and Bloem, 2000), or decided exactly using automata-based language containment checks (Tauriainen, 2006).

Spot implements many (but not all) rewriting rules taken from the aforementioned sources, plus some of its own. A complete listing of all these rules is distributed along with Spot[2] and is too long to be reproduced here. We only discuss a couple of them to

illustrate the point that these rewritings should be selected from the point of view of the translation algorithm that will be used next.

### 4.1   A harmful rewriting rule

As a first example, we do not apply the rule $\mathsf{F}(\varphi \wedge \mathsf{G}\,\mathsf{F}\,\psi) \equiv (\mathsf{F}\,\varphi) \wedge (\mathsf{G}\,\mathsf{F}\,\psi)$ suggested by Somenzi and Bloem (2000). Intuitively, this rule is dubious because $\mathsf{F}(\varphi \wedge \mathsf{G}\,\mathsf{F}(\psi))$ appears less complex to translate. Indeed, translating $\mathsf{F}\,\Phi$ is just a matter of creating an initial state that accepts any letter for a finite number of step, and non-deterministically jumps into a state that will recognise $\Phi$ when a letter matching the beginning of $\Phi$ is found. However, translating a formula such as $(\mathsf{F}\,\varphi) \wedge \mathsf{G}\,\mathsf{F}\,\psi$ is harder because in the initial state you have four choices to consider: either the input can be the start of $\varphi$, or it is the start of $\psi$, or it is both, or it is none. When $\varphi$ and $\psi$ are atomic propositions as in Figure 9, these four cases can be reduced to three. It turns out that on the automaton $\mathcal{A}_2$ from Figure 9 the states $(\mathsf{F}\,a) \wedge \mathsf{G}\,\mathsf{F}\,b$ and $(\mathsf{F}\,a) \wedge (\mathsf{F}\,b) \wedge \mathsf{G}\,\mathsf{F}\,b$ have exactly the same outgoing transitions: they can be merged. Thanks to the BDD identification discussed in Section 3.2, Spot will actually output an automaton similar to $\mathcal{A}_1$ for both formulas $\mathsf{F}(a \wedge \mathsf{G}\,\mathsf{F}\,b)$ or $(\mathsf{F}\,a) \wedge \mathsf{G}\,\mathsf{F}\,b$. This is not the case when $\varphi$ and $\psi$ are more complex.

**Figure 9**   Paper-and-pen translations into TGBA of $\mathsf{F}(a \wedge \mathsf{G}\,\mathsf{F}\,b)$ and $(\mathsf{F}\,a) \wedge \mathsf{G}\,\mathsf{F}\,b$
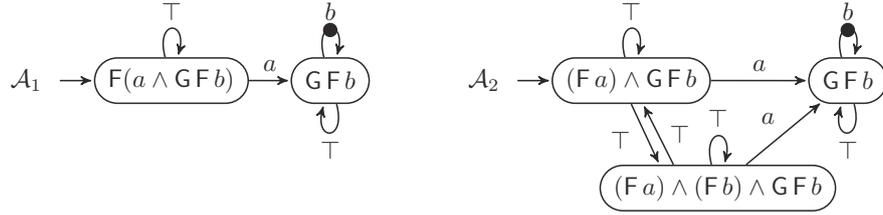


This rewriting rule, which we applied in the past, also prevented other useful rules to apply. E.g., Spot 0.5 would rewrite the formula $\mathsf{F}(\varphi_1 \wedge \mathsf{G}\,\mathsf{F}\,\psi_1) \vee \mathsf{F}(\varphi_2 \wedge \mathsf{G}\,\mathsf{F}\,\psi_2)$ as $((\mathsf{F}\,\varphi_1) \wedge \mathsf{G}\,\mathsf{F}\,\psi_1) \vee (\mathsf{F}\,\varphi_2) \wedge \mathsf{G}\,\mathsf{F}\,\psi_2)$ missing the opportunity to apply the rule $\mathsf{F}(\Psi_1) \vee \mathsf{F}(\Psi_2) = \mathsf{F}(\Psi_1 \vee \Psi_2)$. Since Spot 0.6, we rewrite this formula as $\mathsf{F}((\varphi_1 \wedge \mathsf{G}\,\mathsf{F}\,\psi_1) \vee (\varphi_2 \wedge \mathsf{G}\,\mathsf{F}\,\psi_2))$, which is easier to translate for similar reasons.

### 4.2   Handling the W and M operators

Figure 2 includes rules to translate the W (weak until) and M (strong release) LTL operators. Many tools dealing with LTL formulas do not implement these operators or treat them as syntactic sugar: they do not add expressive power and can be rewritten using other operators. To illustrate the importance of the rewriting rules from the point of view of the translator algorithm, we consider different rewritings for these operators.

The formula $a\,\mathsf{W}\,b$ is usually rewritten into $(a\,\mathsf{U}\,b) \vee \mathsf{G}\,a$. For instance this is the implementation of the W operator in Spin 6.2.2. From the point of view of an LTL translator based on a tableau method, this is not a very good rewriting as it requires a non-deterministic choice between $a\,\mathsf{U}\,b$ and $\mathsf{G}\,a$ at the very beginning. A better

rewriting is $a \, \mathsf{W} \, b \equiv a \, \mathsf{U}(b \vee \mathsf{G} \, a)$, as it postpones the choice between $b$ and $\mathsf{G} \, a$. This latter rewriting was used by Dwyer et al. (1998), although they now changed their web site (http://patterns.projects.cis.ksu.edu/documentation/patterns/ltl.shtml) to use $\mathsf{W}$ for simplicity. An even better choice, although less intuitive, is $a \, \mathsf{W} \, b \equiv b \, \mathsf{R}(a \vee b)$, since no promise have to be introduced.

The TGBAs corresponding to the translation of these different rewritings are shown on Figure 10. Similar rewritings for $\mathsf{M}$ exist:

$$a \, \mathsf{M} \, b \equiv (a \, \mathsf{R} \, b) \wedge (\mathsf{F} \, a) \equiv a \, \mathsf{R}(b \wedge \mathsf{F} \, a) \equiv b \, \mathsf{U}(a \wedge b).$$

**Figure 10** Four formulas equivalent to $a \, \mathsf{W} \, b$ and their corresponding automata
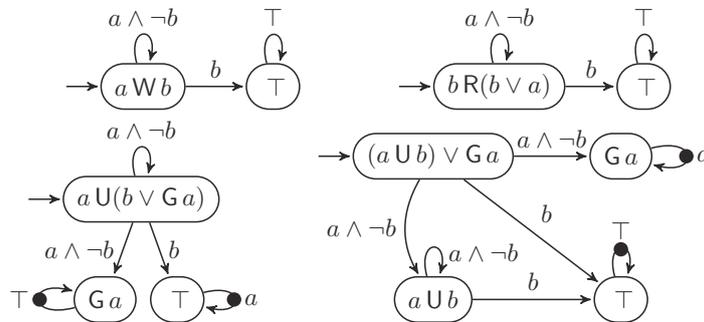


Since Spot fully supports the $\mathsf{W}$ and $\mathsf{M}$ operators, our basic rewriting rules actually perform the reverse of all the previous rewritings (e.g., we rewrite into $a \, \mathsf{W} \, b$ the formulas $(a \, \mathsf{U} \, b) \vee \mathsf{G} \, a$, $a \, \mathsf{U}(b \vee \mathsf{G} \, a)$, and $b \, \mathsf{R}(a \vee b)$).

### 4.3   Implementation of LTL formulas

The implementation of all these rewriting rules benefit greatly from our representation of a set of LTL formulas as a forest of 'syntax DAGs' with sharing of subformulas.

LTL formulas are reference counted and a unicity table makes sure that two equal formulas (or subformulas) will share the same address. The operators $\wedge$ and $\vee$ are handled as $n$-ary operators, and their operands are always sorted. We can therefore easily detect that $a \wedge \mathsf{X}(b) \wedge \mathsf{F}(c)$ is equivalent to $\mathsf{F}(c) \wedge \mathsf{X}(b) \wedge a$ because the two formula objects will have the same address.

The uniqueness of each subformula also helps to speed up rewriting algorithms, as they use a cache when processing subformulas recursively.

## 5   Post-processings

Once an LTL formula has been translated into a TGBA as described in Section 3, Spot implements different kinds of post-processings. We first describe each processing independently before explaining when there are used and how they are chained.

## 5.1  SCC pruning

It may happen that the TGBA constructed by the translation contains states that do not contribute to its language. A classical optimisation is therefore to remove all non-accepting SCCs that cannot reach an accepting SCC (Somenzi and Bloem, 2000).

Since we have to traverse the entire automaton to classify its SCCs as accepting or non-accepting, we can also perform a few other improvements along the way. The acceptance marks of transitions that do not belong to an accepting SCC can be removed. Similarly, acceptance marks that are always present at the same time as another acceptance mark can be simplified.

## 5.2  Minimisation of weak deterministic BA

A Büchi automaton is weak if, in each SCC, either all the cycles are accepting, or all cycles are non-accepting.

It is well known that not all BA can be determinised [Vardi, (1996), Prop. 8]. There is a subclass of properties that can be represented by weak deterministic Büchi automata (WDBA), and for which there exists an algorithm to compute the minimal WDBA recognising the property (Löding, 2001). This class corresponds to the '*obligations*' in the temporal hierarchy of Manna and Pnueli (1990) and includes a large number of LTL formulas used for model checking.For instance 40 formulas out of the 55 formulas from Dwyer et al. (1998) are obligations.

Dax et al. (2007) showed how to implement this minimisation without knowing *a priori* if the translated property actually is an obligation: the correctness of the minimisation is tested *a posteriori* using a language equivalence test (easy to implement because a WDBA can be complemented like deterministic finite automata, and the original TGBA can be complemented by translating the negation of the property).

Dax et al. (2007) did a comparison of the size produced by different translators (not Spot, which they did not know) with the size of the minimal WDBA. This revealed that although it was deterministic, the minimal WDBA usually had a number states smaller or equal to that of the automata produced by the translators.

This WDBA minimisation has since been integrated into Spot, and we completed the benchmark of Dax et al. (2007) in the previous version of this paper (Duret-Lutz, 2011). Our implementation takes a TGBA, and outputs a deterministic Büchi automaton when the WDBA-minimisation is valid. We avoid the language equivalence test in a number of cases by testing whether the translated formula actually belongs to the syntactic obligation class (Černá and Pelánek, 2003).

While being able to output a minimal deterministic automaton for some class of LTL formulas is appreciable, we have found a few cases were using such a deterministic output was not desirable because the deterministic automaton was too large.

As an example consider the family of LTL formulas $\mathsf{G}\,p_1 \vee \mathsf{G}\,p_2 \vee \ldots \mathsf{G}\,p_n$. Figure 11 shows the result of the translation for $n = 3$ before and after WDBA minimisation. The non-deterministic automaton has $n + 1$ states, while the minimal deterministic automaton has $2^n - 1$ states. Experiments on actual model checking problems show that the smaller of these two automata has to be preferred, despite its non-determinism, when the full product with the system must be constructed.

**Figure 11** Left: translation of $\mathsf{G}\,a \vee \mathsf{G}\,b \vee \mathsf{G}\,c$. Right: its minimal WDBA
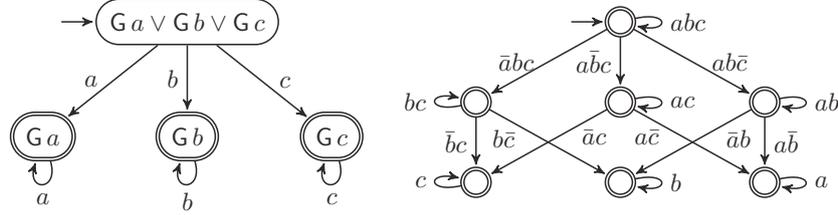


## 5.3 Simulation-based reductions

Spot implements the simulation-based reductions described by Somenzi and Bloem (2000), which are easily adjusted to work on a TGBA. Intuitively direct simulation can merge states based on the inclusion of the sets of infinite runs *starting from* these states, while reverse simulation would merge states based on the inclusion between sets of finite runs *leading to* these states.

Our implementation has the same structure as the *StrongFairSimulation* algorithm of Etessami and Holzmann (2000), except that we represent the class (or colour) of a state using BDD variables to ease inclusion checks. More details about our implementation are given by Babiak et al. (2013).

## 5.4 Degeneralisation

A degeneralisation algorithm takes a generalised automaton with $n$ states and $m$ acceptance marks, and produces a Büchi automaton with at most $n(m+1)$ states. The classical algorithm used to TGBA into BA (Clarke et al., 2000, Section 9.2.2) can be adapted to transform TGBA into Büchi automata (Giannakopoulou and Lerda, 2002; Gastin and Oddoux, 2001) as follows.

If $\mathcal{T} = \langle AP, \mathcal{Q}, q_0, \mathcal{F}, \delta \rangle$ is a TGBA with $m$ acceptance marks $\mathcal{F} = \{f_1, f_2, \ldots, f_m\}$, then an equivalent Büchi automaton $\mathcal{T} = \langle AP, \mathcal{Q}', q_0', \mathcal{F}', \delta' \rangle$ can be constructed as follows:

- $\mathcal{Q}' = \mathcal{Q} \times \{0, \ldots, m\}$, i.e., the original automaton is cloned in $m+1$ levels

- $\mathcal{F}' = \mathcal{Q} \times \{m\}$, i.e., states from the last level are accepting

- $\delta' = \{((s, j), l, (d, level_j(F))) \mid (s, l, F, d) \in \delta\}$
  where $level_j(F) = \begin{cases} 0 & \text{if } j = m \\ j+1 & \text{if } j < m \text{ and } f_{j+1} \in F \\ j & \text{otherwise} \end{cases}$

  i.e., for each level $j < m$ the outgoing transitions that carry $f_{j+1}$ are redirected to the next level and all transitions from the last level are redirected to level $0$

- $q_0' = (q_0, 0)$, i.e., the initial state is on the first level (but any other level would also be correct).

This setup guarantees that any accepting path in the degeneralised automaton will correspond to an infinite path that sees all acceptance marks infinitely often in the original automaton. The classical optimisation is to 'jump levels', i.e., when a transition from level $i < m$ carries acceptance marks $f_{i+1}$, $f_{i+2}$, and $f_{i+3}$, it can be redirected to the level $i + 3$. This corresponds to the following redefinition of $level_j(F)$:

$$level_j(F) = \begin{cases} \max\{n \in \{j, \ldots, m\} \mid \forall k \in \{j + 1, \ldots, n\}, f_k \in F\} & \text{if } j < m, \\ \max\{n \in \{0, \ldots, m\} \mid \forall k \in \{1, \ldots, n\}, f_k \in F\} & \text{if } j = m. \end{cases}$$

The automaton $\mathcal{B}_1$ from Figure 1 was degeneralised from $\mathcal{T}_1$ with this definition, in the order $f_1 = \bigcirc$, $f_2 = \bullet$, and setting the initial state in the last level.

Another optimisation this is implemented in Spot is a *pulling of acceptance marks*. When all outgoing transitions of a state $s$ have a set $Y$ of acceptance marks in common, this set can be added to the acceptance marks of all the incoming transitions. This is correct because if a run traverses $s$ it will necessarily see all acceptance marks from $Y$; it makes no difference if its sees them twice.

This degeneralisation procedure offers $m!$ possible ways to order the acceptance marks, and there are $m + 1$ possible levels on which the initial state can be located. Changing these parameters might make some states from $\mathcal{Q} \times \{0, \ldots, m\}$ unreachable, and can thus reduce the automaton. For one TGBA, we therefore have $m!(m + 1)$ possible degeneralisations using only this definition.

In Spot, the order of acceptance of sets used for the degeneralisation correspond to the order in which the corresponding promises where introduced during the translation, and the initial state is always on the first level. There is definitely room for improvement here, since the initial submission of this paper, we have been working with the authors of `ltl3ba` to improve the situation (Babiak et al., 2013).

Oddoux (2003, Section 6.1.2) mentions another kind of degeneralisation in which the acceptance marks can be taken in any order and where each state of the degeneralised automaton has to retain the set of all acceptance marks that are waited for. This can potentially multiply the size of the original automaton with $2^m$ if $m$ acceptance marks are used. But this might be worth a try when $m$ is very small.

## 5.5  *The complete post-processing chain*

Because it is not always clear in which context the translated automaton will be used, Spot 1.0 introduces two different options to specify the intent of the translation.

- `deterministic` is used to indicate that an output that is (as much as possible) deterministic is desired. E.g., the right automaton of Figure 11 should be preferred.

  In this case, we first prune useless SCC and acceptance marks in the translated TGBA, then we apply WDBA-minimisation. If the latter succeeded, we output its result (a Büchi automaton) as-is. In case where WDBA-minimisation was not applicable, we reduce the TGBA by iterating both direct and reverse simulation until the automaton is not reduced any more. The simulation-reduced TGBA is then degeneralised if requested.

- `small` is used to indicate that an output with less states should be favoured. We shall still strive to make it deterministic, but if a choice like that of Figure 11 happens, we will prefer the left automaton.

  The post-processing for this intent also starts by pruning useless SCCs and acceptance marks. Then we compute two different automata, and return the smallest: the first automaton is the result of WDBA-minimisation (if that result exists), and the second is the result of the iterated simulation (optionally degeneralised). If the two automata have an equal number of states, we keep the WDBA because it is guaranteed to be deterministic.

## 6 Benchmarks

The following sections present different benchmarks comparing Spot with other translators that are publicly available (including older versions of Spot).
  These translators (presented in chronological order) are:

- The Spin model checker. Its `-f` option converts an LTL formula into a never claim representing a (degeneralised) Büchi automaton. Spin's LTL translator is based on the tableau construction of Gerth et al. (1996). Spin has some trivial and unconditional rewriting rules for LTL, and includes simple post-processings.

- LBT (Rönkkö, 1999) also implements the translation of Gerth et al. (1996), but produces a generalised Büchi automaton. LBT only apply trivial rewriting rules. It has no post-processings.

- `wring` (Somenzi and Bloem, 2000) implements some unconditional LTL rewritings, as well as some implication-based checks. Using a tableau construction it builds a generalised Büchi automaton with labels on states (rather than transitions). This GBA is simplified using SCC-based and simulation-based reductions.

- `ltl2ba` (Gastin and Oddoux, 2001) is a descendant of Spin's translator in the sense that it reused the same code base. However the translation algorithm has been completely rewritten. LTL formulas are reduced using all classes of rewriting rules (the implication checks are syntactic), translated into an intermediate alternating Büchi automaton, which is then converted into a TGBA, which is finally degeneralised into a Büchi automaton. Some simplifications (like removing redundant transitions and useless SCCS) are performed at all these steps.

- `modella` (Sebastiani and Tonetta, 2003) uses a tableau construction, implements all classes of rewriting rules (with syntactic implication checks), it also implements simulation-based reductions on the Büchi automaton. One of the main points of Modella's authors was that it is worth improving the determinism of the automaton at the expense of its size, because this will pay off when this automaton is later synchronised with a system to check.

- `ltl2nba` (Fritz, 2003) translates LTL formulas into alternating BA with $\varepsilon$-transitions, and performs simulation reductions directly on that. These alternating automata are then converted into Büchi automaton using the Miyano-Hayashi construction. No pre- or post- processing are performed.

- `ltl3ba` (Babiak et al., 2012) is a reimplementation of `ltl2ba` in C++ with better data structures and additional optimisations. It implements many LTL rewriting rules, including some new ones based on a class for formulas called alternating formulas. For instance it uses BDDs to simplify the guards of transitions. It implements a technique called suspension that would effectively solve the problem described in 4.1: when `ltl3ba` translates $(\mathsf{F}\,\varphi) \wedge (\mathsf{G}\,\mathsf{F}\,\psi)$, it suspends the translation of $\mathsf{G}\,\mathsf{F}\,\psi$ until a point where $\varphi$ must hold. This translator also implements a direct-simulation reduction on the final Büchi automaton (option `-S`), and has an option to improve determinism (option `-M`).

`ltl2nba` and `wring` are scripts written respectively in Python and Perl. All other tools are compiled from C or C++. All the following experiments were ran under GNU/Linux on an Intel Core2 Q9550 running at 2.83GHz with 8GB of RAM.

### 6.1   184 LTL formulas from the literature

The benchmark consists in 92 LTL formulas:

- 55 formulas from Dwyer et al. (1998) (where $a\,\mathsf{W}\,b$ was written as $a\,\mathsf{U}(b \vee \mathsf{G}\,a)$)

- 25 formulas from Somenzi and Bloem (2000) – their paper shows 27 formulas but two of them are already the negations of other formulas in the list

- 12 formulas from Etessami and Holzmann (2000).

With their negations this makes a total of 184 formulas.

A summary of the translation of these formulas is presented in Table 1. The statistics displayed in this table were gathered using `ltlcross`, a Spot-based reimplementation of LBTT (Tauriainen and Heljanko, 2002) that cross compares translators in order to detect errors (for our extensive test suite) and collect statistics (for our papers). The tools have been clustered by type of automaton produced, with Spot appearing in two groups depending on whether it was configured to output BA or TGBA.

As shown by the *count* column, Spin failed to translate 11 formulas within the ten minutes limit we had set up (the machine was swapping before the end of these ten minutes, meaning spin needed more than the available memory). Wring aborted in three cases with an error message from Perl. Modella produced one incorrect automaton.

Modella and `ltl2nba` output automata in a format in which states need not be declared accepting if they all are, this explains why they show less acceptance sets/marks than translated formulas (but this difference is not important). For other LTL-to-BA translators we used the never claim output.

**Table 1** Cumulative summary for the translation of 184 formulas by each tool

| Translator | Count | Automata sizes | | | | | Non-det. | | Time | Products | | |
| | | States | Edges | Trans. | Acc. | SCC | States | Aut. | | States | Trans. | SCC |
| --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- | --- |
| **BA** | | | | | | | | | | | | |
| spin 6.2.2 | 173 | 2,166 | 17,428 | 62,771 | 173 | 1,242 | 1,861 | 169 | 964.09 | 287,823 | 9,983,437 | 29,575 |
| lt12ba 1.1 | 184 | 1,017 | 3,385 | 30,237 | 184 | 732 | 811 | 173 | 0.67 | 194,684 | 5,638,095 | 4,390 |
| lt12nba | 184 | 952 | 3,065 | 27,158 | 181 | 724 | 744 | 174 | 18.77 | 179,988 | 5,472,936 | 3,470 |
| modella 1.5.9 | 183 | 1,312 | 4,391 | 23,998 | 180 | 861 | 650 | 119 | 30.78 | 219,548 | 4,258,427 | 14,753 |
| lt13ba 1.0.1 −S | 184 | 812 | 2,242 | 21,339 | 184 | 627 | 607 | 172 | 1.00 | 154,567 | 4,350,636 | 1,769 |
| lt13ba 1.0.1 −M | 184 | 875 | 2,634 | 15,101 | 184 | 666 | 357 | 124 | 1.05 | 163,669 | 3,090,518 | 4,133 |
| lt13ba 1.0.1 −M −S | 184 | 848 | 2,437 | 14,495 | 184 | 648 | 330 | 124 | 1.07 | 158,517 | 2,953,167 | 3,260 |
| Spot 1.0 deterministic | 184 | 683 | 1,707 | 10,627 | 184 | 496 | 93 | 45 | 4.27 | 132,239 | 2,409,181 | 1,453 |
| Spot 1.0 small | 184 | 678 | 1,683 | 10,517 | 184 | 495 | 99 | 50 | 4.47 | 131,850 | 2,406,921 | 1,474 |
| **GBA** | | | | | | | | | | | | |
| lbt 1.2.2 | 184 | 8415 | 130,501 | 884,155 | 333 | 4,781 | 7,588 | 180 | 1.62 | 1,423,870 | 108,033,745 | 596,754 |
| Wring 1.1.0 | 181 | 1,476 | 4,943 | 49,416 | 182 | 1,120 | 1,027 | 166 | 31.80 | 253,354 | 7,765,883 | 74,934 |
| **TGBA** | | | | | | | | | | | | |
| Spot 1.0 deterministic | 184 | 641 | 1,573 | 9,964 | 198 | 493 | 84 | 45 | 4.41 | 124053 | 2,258,314 | 1,446 |
| Spot 1.0 small | 184 | 636 | 1,549 | 9,854 | 198 | 492 | 90 | 50 | 4.56 | 123,664 | 2,256,054 | 1,467 |

Notes: With the exception of the *count* column, smaller values are better. The first data column displays the total *count* of formulas successfully translated by each tool. The remaining columns all display accumulated values over all successful translations. The first columns display the total number of *states, edges, transitions, acceptance marks* (1 for BA), and *strongly connected components* in all the translated automata. They are followed by the total number of *non-deterministic states*, and the total number of *non-deterministic automata* in these translated automata. The next column gives number of seconds it took to translate all formulas. For each formula, one random (and deadlock free) state space of 200 states was created, and used to builds a synchronous product with each translated automaton. The size of this product is shown in the last three columns. We do not have to distinguish edges from transitions in the product because all atomic propositions are valued in the state space: each edge maps to one transition.

The product with a random state space gives some idea of the behaviour of the automaton during model checking. The intuition is that if this state space was that of a real model to verify, the model checking procedure would need space proportional to the number of state in the product, and time proportional to the number of transitions in that product. This can be used to argue for instance that although `modella`'s automata are bigger than `ltl2ba`'s, they will yield less transition in the product, and therefore make a faster verification. A similar effect can be seen with `ltl3ba`'s `-M` option: it improves the determinism at the expense of the number of states, but this pays off in terms of transitions in the products. This interpretation of the last columns should be mitigated by the fact that these random state spaces are not real models (Pelánek, 2008), and that an actual model checker will implement other techniques to avoid constructing the entire product.

The statistics for Wring are slightly unfair because the nature of the automata it generates (state-based labels, and multiple initial states) is very different from the automata produced by other translators. In order to integrate Wring in our benchmark, we had to add a fake initial state (connected to all the former initial states) to each produced automata, and move the label of each state onto all its incoming transitions. This quick transformation adds one more state per automaton, and is also accounted for in the total run time.

On Table 1 Spot appears better than all other translator on all accounts except its run time (which is still reasonable). The number of states in its BA output is probably even more impressive if we additionally consider the number of nondeterministic automata: the automata are smaller *and* more deterministic.

However these cumulative values hide the actual differences that may be observed when comparing the results formula by formula. There is only one formula for which Spot produces an automaton with one more states than other tools. However if we distinguish automata with equal number of states by their number of transitions, then these cases are more numerous as shown in Table 2.

**Table 2**    Comparing LTL-to-BA translator on 182 formulas

|  | spin | ltl2ba | ltl2nba | modella | ltl3ba -S | ltl3ba -M | ltl3ba -M -S | Spot (det.) | Spot (small) |
|---|---|---|---|---|---|---|---|---|---|
| spin |  | 131 | 132 | 134 | 135 | 166 | 168 | 169 | 169 |
| ltl2ba | 2 |  | 38 | 69 | 47 | 160 | 162 | 156 | 159 |
| ltl2nba | 5 | 25 |  | 65 | 43 | 147 | 150 | 158 | 161 |
| modella | 35 | 104 | 108 |  | 116 | 116 | 116 | 132 | 132 |
| ltl3ba -S | 0 | 18 | 31 | 55 |  | 138 | 144 | 153 | 157 |
| ltl3ba -M | 0 | 3 | 18 | 4 | 22 |  | 19 | 82 | 84 |
| ltl3ba -M -S | 0 | 3 | 16 | 4 | 15 | 0 |  | 79 | 82 |
| Spot (det.) | 0 | 10 | 9 | 1 | 10 | 24 | 25 |  | 5 |
| Spot (small) | 0 | 7 | 6 | 1 | 6 | 20 | 20 | 0 |  |

Notes: The value on line $i$ and column $j$ shows how many times the automaton produced by translator #$i$ was strictly bigger than the automaton produced by translator #$j$.

Here 'bigger' means 'more states' or 'equal number of states and more transitions'.

Detailed results can be found at http://www.lrde.epita.fr/ adl/ijccbs/, and include an interactive page that can build Table 2 for different comparison criteria.

### 6.2 Some formulas for which the minimal Büchi automaton is known

Cichoń et al. studied several classes of LTL formulas for which they calculated the size (in states) of the minimal Büchi automaton that could represent the property. They compared the output of Spot 0.4 and `ltl2ba` 1.1, neither of which was able to translate all formulas efficiently. Sometimes they would take too long (hours or days), sometimes they would produce automata larger than necessary.

Here are the five families of formulas they evaluated on both tools for $n$ ranging from 1 to 20:

$$\alpha_n = \mathsf{F}(p_1 \wedge \mathsf{F}(p_2 \wedge \mathsf{F}(\ldots \mathsf{F}\,p_n))) \wedge \mathsf{F}(q_1 \wedge \mathsf{F}(q_2 \wedge \mathsf{F}(\ldots \mathsf{F}\,q_n)))$$

$$\beta_n = \mathsf{F}(\underbrace{p \wedge \mathsf{X}(p \wedge \mathsf{X}(p \wedge \ldots))}_{n \text{ occurrences of } p}) \wedge \mathsf{F}(\underbrace{q \wedge \mathsf{X}(q \wedge \mathsf{X}(q \wedge \ldots))}_{n \text{ occurrences of } q})$$

$$\beta'_n = \mathsf{F}(\underbrace{p \wedge \mathsf{X}(p) \wedge \mathsf{X}\,\mathsf{X}(p) \wedge \ldots}_{n \text{ occurrences of } p}) \wedge \mathsf{F}(\underbrace{q \wedge \mathsf{X}(q) \wedge \mathsf{X}\,\mathsf{X}(q) \wedge \ldots}_{n \text{ occurrences of } q})$$

$$\varphi_n = \mathsf{G}\,\mathsf{F}\,p_1 \wedge \mathsf{G}\,\mathsf{F}\,p_2 \wedge \ldots \wedge \mathsf{G}\,\mathsf{F}\,p_n$$

$$\psi_n = \mathsf{F}\,\mathsf{G}\,p_1 \vee \mathsf{F}\,\mathsf{G}\,p_2 \vee \ldots \vee \mathsf{F}\,\mathsf{G}\,p_n$$

Nowadays Spot, as well as the new `ltl3ba`, are both able to translate all 100 formulas into their optimal BA, and within reasonable time. Table 3 shows the evolution of the total time required to translate the 100 formulas.

**Table 3** Total time required to translate $\alpha_n$, $\beta_n$, $\beta'_n$, $\varphi_n$, and $\psi_n$ for $1 \leq n \leq 20$

| | |
|---|---|
| Spot 0.8.3 | 562 seconds |
| Spot 0.9 | 315 seconds |
| Spot 0.9.1 | 198 seconds |
| Spot 1.0 | 150 seconds |
| ltl3ba 1.0.1 | 77 seconds |

The translation of $\alpha_n$, $\beta_n$, $\beta'_n$ and $\psi_n$ is nearly instantaneous: $\alpha_{20}$, the longest to translate, requires 1.7s. Therefore the larger part of the run time summed in Table 3 comes from the $\varphi_n$ family of formulas. For instance Spot 1.0 spends 90 s computing just $\varphi_{20}$, and 40% of that time is spent in our inefficient degeneralisation procedure. Comparatively, `ltl3ba`, which has some better handling for subformulas of that form, will translate $\varphi_{20}$ in only 42 s.

The variation between Spot 0.9 and Spot 0.9.1 corresponds to the introduction of the optimised translation of $\mathsf{G}$ discussed in Section 3.4, which is especially pertinent for the translation of $\varphi_n$: our LTL pre-processing will rewrite $\varphi_n$ as $\mathsf{G}(\mathsf{F}(p1) \wedge \mathsf{F}(p2) \wedge \ldots \wedge \mathsf{F}(p_n))$ and from there the rules from Figure 7 will avoid the creation of many useless $\mathsf{Nxt}[]$ variables for each of the $\mathsf{F}$ subformulas.

## 6.3   LTL counter

Rozier and Vardi (2007) compared 9 LTL translators, on various families of LTL formulas.

The first family of formulas they experimented on is scalable. For a given $n$ they generated an LTL formula $C_n$ that matches an infinite sequence of bits in which all the values of a $n$-bit counter have been concatenated. E.g., $C_3 = ((a \wedge (\mathsf{G}(a \rightarrow (\mathsf{X}(\neg a \wedge \mathsf{X}(\neg a \wedge \mathsf{X} a)))))) \wedge ((\neg b) \wedge \mathsf{X}(\neg b \wedge \mathsf{X} \neg b)) \wedge (\mathsf{G}((a \wedge \neg b) \rightarrow (\mathsf{X}((\mathsf{X} \mathsf{X} b) \wedge (((\neg a) \wedge (b \rightarrow \mathsf{X} \mathsf{X} \mathsf{X} b) \wedge ((\neg b) \rightarrow (\mathsf{X} \mathsf{X} \mathsf{X} \neg b))) \mathsf{U} a))))) \wedge (\mathsf{G}((a \wedge b) \rightarrow (\mathsf{X}((\mathsf{X} \mathsf{X} \neg b) \wedge ((b \wedge (\neg a) \wedge \mathsf{X} \mathsf{X} \mathsf{X} \neg b) \mathsf{U}(a \vee ((\neg a) \wedge (\neg b) \wedge (\mathsf{X}((\mathsf{X} \mathsf{X} b) \wedge (((\neg a) \wedge (b \rightarrow \mathsf{X} \mathsf{X} \mathsf{X} b) \wedge ((\neg b) \rightarrow \mathsf{X} \mathsf{X} \mathsf{X} \neg b)) \mathsf{U} a))))))))))))).$[3] Such a formula will match a sequence consisting of $\begin{smallmatrix} a:\ 1\ 0\ 0\ 1\ 0\ 0\ 1\ 0\ 0\ 1\ 0\ 0\ 1\ 0\ 0\ 1\ 0\ 0\ 1\ 0\ 0\ 1\ 0\ 0 \\ b:\ 0\ 0\ 0\ 1\ 0\ 0\ 0\ 1\ 0\ 1\ 1\ 0\ 0\ 0\ 1\ 1\ 0\ 1\ 0\ 1\ 1\ 1\ 1\ 1 \end{smallmatrix} \cdots$ repeated infinitely. Variable $a$ signals the start of each value, while variable $b$ iterates over the three bits of each value from least to most significant bit (000, 100, 010, …).

From this description it should be clear that the smallest automaton that can recognise $C_n$ is a deterministic loop with $n2^n$ states and as many transitions. Figure 12 shows this automaton for $C_3$. Any translator that constructs such an automaton explicitly will have a run time that is worse than exponential in $n$.

**Figure 12**   A Büchi automaton that recognises $C_3$



**Figure 13**   Run time of different tools on the translation of LTL counter formulas (see online version for colours)



Notes: Spot 0.4 was the version used by Rozier and Vardi (2007) in their experiments. Spot 0.7 was the version used for our experiments at VECOS '11 (Duret-Lutz, 2011)

Figure 13 shows the run time taken by `ltl2ba`, `ltl3ba`, and three versions of Spot to translate $C_n$ for increasing $n$. Executions were limited to 15min. Other tools are not shown, as they already fail to translate $C_4$ (sometimes even $C_1$) within this limit.

All these tools have been configured to not perform any pre- and post-processings. Also we patched `ltl2ba`, `ltl3ba`, so they would stop right after having constructing a TGBA without constructing a Büchi automaton. Therefore we are only measuring the scalability of the core LTL-to-TGBA translation algorithm in each of these tools. (We verified that each tool was slower with pre-processing turned on, which means that LTL rewriting are of no help on the $C_n$ family of formulas.)

## 7  Conclusions

We have presented the main ingredients of the translation module of Spot. The core of the translation is the BDD-based tableau construction of Couvreur (1999), which has been extended in several ways: more determinism (a suggestion of Couvreur himself), some simplifications of the translations of subformulas of G, and a reduction of the number of promises required. This translation is preceded by a huge number of LTL rewriting rules, and followed by several post-processings to reduce the size of the produced automaton. This entire chain produces small automata that tend to be very deterministic, although not necessarily as fast as other translators such as `ltl3ba`.

Our implementation is extensively tested using both handwritten and random LTL formulas, and cross-compared to other translators using tools such LBTT (Tauriainen and Heljanko, 2002) or `ltlcross`, a Spot-based reimplementation.

The degeneralisation algorithm appears to be a weak point in Spot, and the historical reason is that we seldom use it in practice. When building a model checker on top of Spot, we implement the automata-theoretic approach using TGBA directly.

Although we have not discussed this, the implementation of this translation in Spot has been extended to support PSL (Accellera, 2004), and our post-processings also include algorithms to output monitors (Tabakov and Vardi, 2010) and testing automata (BenSalem et al., 2011, 2012).

It has been argued (Cichoń et al.; Tsay et al., 2011) that rather than optimising an algorithm to try to produce the best automata always, it would be useful to create a database of optimal automata for commonly used formulas. However different uses may call for different definition of *optimal automaton*. In the context of model-checking, one usually wants to reduce the size of the product of the property with the system, and translating the property into a small automaton that is the most deterministic possible usually helps (Sebastiani and Tonetta, 2003), but it is not always clear if more determinism justify additional states. In the context of monitoring, where an automaton is monitoring a running process, a deterministic automaton is preferred. In the context of synthesis of reactive systems, Ehlers and Finkbeiner (2010) prefers to minimise the number of states at the expense of determinism.

Also different kinds of automata can be used for verification: model checking with TGBA is usually better than model checking with Büchi automata when the formula incur a lot of acceptance marks (Couvreur et al., 2005). Using testing automata also appears promising (Geldenhuys and Hansen, 2006; BenSalem et al., 2011). A database should therefore not be limited to BA.

While we agree that such a database, like the *Büchi Store* project (Tsay et al., 2011), is useful, we still believe that it is important to have a translation that is efficient and versatile enough to be tuned to the needs of a particular situation.

## Acknowledgements

## References

Accellera (2004) *Property Specification Language Reference Manual v1.1* [online] http://www.eda.org/vfv/ (accessed September 2013).

Babiak, T., Křetínský, M., Řeehák, V. and Strejček, J. (2012) 'LTL to Büchi automata translation: fast and more deterministic', in *Proc. of the 18th Int. Conf. on Tools and Algorithms for the Construction and Analysis of Systems (TACAS '12)*, Vol. 7214 of *LNCS*, pp.95–109, Springer.

Babiak, T., Badie, T., Duret-Lutz, A., Křetínský, M. and Strejček, J. (2013) 'Compositional approach to suspension and other improvements to LTL translation', in *Proceedings of the 20th International SPIN Symposium on Model Checking of Software (SPIN '13)*, Vol. 7976 of *Lecture Notes in Computer Science*, pp.81–98, Springer.

BenSalem, A.E., Duret-Lutz, A. and Kordon, F. (2011) 'Generalized Büchi automata versus testing automata for model checking', in *Proc. of the second International Workshop on Scalable and Usable Model Checking for Petri Net and other models of Concurrency (SUMO '11)*, Vol. 626 of *Workshop Proceedings*, CEUR, Newcastle, UK.

BenSalem, A.E., Duret-Lutz, A. and Kordon, F. (2012) 'Model checking using generalized testing automata', in *Transactions on Petri Nets and Other Models of Concurrency (ToPNoC VI)*, Vol. 7400 of *LNCS*, pp.94–112, Springer.

Bryant, R.E. (1986) 'Graph-based algorithms for Boolean function manipulation', *IEEE Transactions on Computers*, Vol. 35, No. 8, pp.677–691.

Černá, I. and Pelánek, R. (2003) 'Relating hierarchy of temporal properties to model checking', in *Proc. of the 28th Int. Symp. on Mathematical Foundations of Computer Science (MFCS '03)*, Vol. 2747 of *LNCS*, pp.318–327, Springer-Verlag, Bratislava, Slovak Republic.

Cichoń, J., Czubak, A. and Jasiński, A. (2009) 'Minimal Büchi automata for certain classes of LTL formulas', in *Proc. of the Fourth Int. Conf. on Dependability of Computer Systems (DEPCOS '09)*, pp.17–24, IEEE Computer Society.

Clarke, E.M., Grumberg, O. and Peled, D.A. (2000) *Model Checking*, MIT Press, Cambridge, MA, USA.

Couvreur, J-M., Duret-Lutz, A. and Poitrenaud, D. (2005) 'On-the-fly emptiness checks for generalized Büchi automata', in *Proc. of the 12th Int. SPIN Workshop on Model Checking of Software (SPIN '05)*, Vol. 3639 of *LNCS*, pp.143–158, Springer.

Couvreur, J-M. (1999) 'On-the-fly verification of temporal logic', in *Proc. of the World Congress on Formal Methods in the Development of Computing Systems (FM '99)*, Vol. 1708 of *LNCS*, pp.253–271, Springer-Verlag, Toulouse, France.

Dax, C., Eisinger, J. and Klaedtke, F. (2007) 'Mechanizing the powerset construction for restricted classes of $\omega$-automata', in *Proc. of the 5th Int. Symp. on Automated Technology for Verification and Analysis (ATVA '07)*, Vol. 4762 of *LNCS*, Springer.

Duret-Lutz, A. and Poitrenaud, D. (2004) 'SPOT: an extensible model checking library using transition-based generalized Büchi automata', in *Proc. of the 12th Int. Symp. on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS '04)*, pp.76–83, IEEE Computer Society Press, Volendam, The Netherlands.

Duret-Lutz, A. (2011) 'LTL translation improvements in Spot', in *Proc. of the 5th Int. Workshop on Verification and Evaluation of Computer and Communication Systems (VECoS '11), Electronic Workshops in Computing*, British Computer Society, Tunis, Tunisia.

Dwyer, M.B., Avrunin, G.S. and Corbett, J.C. (1998) 'Property specification patterns for finite-state verification', in *Proc. of the 2nd Workshop on Formal Methods in Software Practice (FMSP '98)*, pp.7–15, ACM Press, New York.

Ehlers, R. and Finkbeiner, B. (2010) 'On the virtue of patience: minimizing Büchi automata', in *Proc. of the 17th International SPIN Conference on Model Checking Software (SPIN '10)*, Vol. 6349 of *LNCS*, pp.129–145, Springer.

Etessami, K. and Holzmann, G.J. (2000) 'Optimizing Büchi automata', in C. Palamidessi (Ed.): |it Proc. of the 11th Int. Conf. on Concurrency Theory (Concur '00), Vol. 1877 of *LNCS*, pp.153–167, Springer-Verlag, Pennsylvania, USA.

Fritz, C. (2003) 'Constructing Büchi automata from linear temporal logic using simulation relations for alternating Büchi automata', in *Proc. of the 8th Int. Conf. on Implementation and Application of Automata (CIAA '03)*, Vol. 2759 of *LNCS*, pp.35–48, Springer-Verlag, Santa Barbara, California.

Gastin, P. and Oddoux, D. (2001) 'Fast LTL to Büchi automata translation', in *Proc. of the 13th Int. Conf. on Computer Aided Verification (CAV '01)*, Vol. 2102 of *LNCS*, pp.53–65, Springer-Verlag, Paris, France.

Geldenhuys, J. and Hansen, H. (2006) 'Larger automata and less work for LTL model checking', in *Proc. of the 13th Int. SPIN Workshop (SPIN '06)*, Vol. 3925 of *LNCS*, pp.53–70, Springer.

Gerth, R., Peled, D., Vardi, M.Y. and Wolper, P. (1996) 'Simple on-the-fly automatic verification of linear temporal logic', in *Proc. of the 15th Workshop on Protocol Specification Testing and Verification (PSTV '95)*, pp.3–18, Chapman & Hall, Warsaw, Poland.

Giannakopoulou, D. and Lerda, F. (2002) 'From states to transitions: improving translation of LTL formulæ to Büchi automata', in *Proc. of the 22nd IFIP WG 6.1 Int. Conf. on Formal Techniques for Networked and Distributed Systems (FORTE '02)*, Vol. 2529 of *LNCS*, pp.308–326, Springer-Verlag, Houston, Texas.

Löding, C. (2001) 'Efficient minimization of deterministic weak $\omega$-automata', *Information Processing Letters*, Vol. 79, No. 3, pp.105–109.

Manna, Z. and Pnueli, A. (1990) 'A hierarchy of temporal properties', in *Proc. of the Sixth Annual ACM Symposium on Principles of Distributed Computing (PODC '90)*, pp.377–410, ACM, New York, NY, USA.

Minato, S. (1992) 'Fast generation of irredundant sum-of-products forms from binary decision diagrams', in *Proc. of the third Synthesis and Simulation and Meeting International Interchange workshop (SASIMI '92)*, pp.64–73, Kobe, Japan.

Oddoux, D. (2003) *Utilisation des automates alternants pour un model-checking efficace des logiques temporelles linéaires*, PhD thesis, Universitée Paris 7, Paris, France.

Pelánek, R. (2008) 'Properties of state spaces and their applications', *STTT*, Vol. 10, No. 5, pp.443–454.

Rönkkö, M. (1999) 'LBT: LTL to Büchi conversion' [online] http://www.tcs.hut.fi/Software/maria/tools/lbt/, Implements the tableau construction from Gerth et al. (1996) (accessed September 2013).

Rozier, K.Y. and Vardi, M.Y. (2007) 'LTL satisfiability checking', in *Proc. of the 12th Int. SPIN Workshop on Model Checking of Software (SPIN '07)*, Vol. 4595 of *LNCS*, pp.149–167, Springer.

Sebastiani, R. and Tonetta, S. (2003) "More deterministic' vs. 'smaller' Büchi automata for efficient LTL model checking', in *Proc. of the 12th Advanced Research Working Conference on Correct Hardware Design and Verification Methods (CHARME '03)*, Vol. 2860 of *LNCS*, pp.126–140, Springer-Verlag, L'Aquila, Italy.

Somenzi, F. and Bloem, R. (2000) 'Efficient Büchi automata for LTL formulæ', in *Proc. of the 12th Int. Conf. on Computer Aided Verification (CAV '00)*, Vol. 1855 of *LNCS*, pp.247–263, Springer-Verlag, Chicago, Illinois, USA.

Tabakov, D. and Vardi, M.Y. (2010) 'Optimized temporal monitors for SystemC', in *Proc. of the 1st Int. Conf. on Runtime Verification (RV '10)*, Vol. 6418 of *LNCS*, pp.436–451, Springer.

Tauriainen, H. and Heljanko, K. (2002) 'Testing LTL formula translation into Büchi automata', *International Journal on Software Tools for Technology Transfer*, Vol. 4, No. 1, pp.57–70.

Tauriainen, H. (2006) *Automata and Linear Temporal Logic: Translation with Transition-based Acceptance*, PhD thesis, Helsinki University of Technology, Espoo, Finland.

Tsay, Y-K., Tsai, M-H., Chang, J-S. and Chang, Y-W. (2011) 'Büchi store: an open repository of büchi automata', in *Proc. of the 17th Int. Conf. Tools and Algorithms for the Construction and Analysis of Systems (TACAS '11)*, Vol. 6605 of *LNCS*, pp.262–266, Springer.

Vardi, M.Y. (1996) 'An automata-theoretic approach to linear temporal logic', in F. Moller and G.M. Birtwistle (Eds.): *Proc. of the 8th Banff Higher Order Workshop (Banff '94)*, Vol. 1043 of *LNCS*, pp.238–266, Springer-Verlag, Banff, Alberta, Canada.

Vardi, M.Y. (2007) 'Automata-theoretic model checking revisited', in *Proc. of the 8th Int. Conf. on Verification, Model Checking and Abstract Interpretation (VMCAI '07)*, Vol. 4349 of *LNCS*, Springer, Nice, France. Invited paper.

## Notes

1    Promises should not be mistaken for co-Büchi acceptance conditions. A co-Büchi acceptance condition $\mathcal{F}$ accepts runs that stay in $\mathcal{F}$ continuously; conversely a promise accepts runs that do not make the promise continuously.

2    See the file `doc/tl/tl.pdf` in the Spot distribution.

3    After installing Spot, this formula can be generated with `genltl --rv-counter-linear=3`.