## THÈSE DE DOCTORAT DE l'UNIVERSITÉ PIERRE ET MARIE CURIE

Spécialité

Informatique

École doctorale Informatique, Télécommunications et Électronique (Paris)

Présentée par

## Ala Eddine BEN SALEM

Pour obtenir le grade de DOCTEUR de l'UNIVERSITÉ PIERRE ET MARIE CURIE

## Improving the Model Checking of Stutter-Invariant LTL Properties

Soutenue le 25 septembre 2014 devant le jury composé de :

M. Radu MATEESCU, Directeur de recherche à l'INRIA Grenoble
M. Stefan SCHWOON, Maître de conférences à l'ENS Cachan
Mme Béatrice BÉRARD, Professeur à l'UPMC (Paris 6)
M. Didier BUCHS, Professeur à l'Université de Genève
M. Alexandre DURET-LUTZ, Maître de conférences à l'EPITA
M. Fabrice KORDON, Professeur à l'UPMC (Paris 6)

Rapporteur Rapporteur Examinateur Examinateur Encadrant Directeur de thèse











## Acknowledgments

Firstly, I would like to express my sincerest gratitude to my supervisors Fabrice KORDON and Alexandre DURET-LUTZ, for their continuous support and guidance during my PhD work, and for their patience and encouragement that greatly helped me to realize this thesis work.

I also would like to express my warm thanks to the jury members: Radu MATEESCU and Stefan SCHWOON for accepting to review my thesis, Béatrice BÉRARD and Didier BUCHS for their interest in my work.

I am especially grateful to Joël COURTOIS (director of EPITA) and Olivier RICOU (director of LRDE) for giving me the opportunity to pursue my thesis work in the EPITA Research and Development Laboratory (LRDE). I am also grateful to all my colleagues at the LRDE for the excellent and productive working environment, especially, Akim DEMAILLE, Clément DÉMOULINS, Daniela BECKER, Didier VERNA, Edwin CARLINET, Etienne RENAULT, Guillaume LAZZARA, Jonathan FABRIZIO, Myriam ROBERT-SEIDOWSKY, Nicolas BOURTRY, Reda DEHAK, Roland LEVILLAIN, Stefania CALARASANU, Thierry GÉRAUD and Yongchao XU.

Many thanks go to members of MoVe team (LIP6), especially, Denis POITRENAUD, Béatrice BÉRARD, Yann THIERRY-MIEG and Soheib BAARIR for their advices and constructive suggestions, and Etienne RENAULT, Laure MILLET, Matthieu SASSOLAS, Maximilien COLANGE, Raveca OARGA, Yann BEN MAÏSSA and Yan ZHANG for their friendship and assistance.

Last but certainly not least, I would like to thank deeply all of my family for their infinite support during the difficult time of this thesis and throughout my studies, and for all the good times we spent together.

# Contents

Intr	oductio	n	1
1.1	Conte	xt	1
	1.1.1	Model Checking and Automata-theoretic Approach	1
	1.1.2	The State-space Explosion Problem	2
	1.1.3	Explicit versus Symbolic model checking	3
	1.1.4	LTL and Stutter-invariance	3
1.2	Scope	of the Thesis	4
	1.2.1	Existing Work	4
	1.2.2	Contributions	5
Pre	limina	ries	9
The	Traditi	ional Approaches to LTL Model Checking	11
2.1	Introd	uction	11
2.2	Model	ling the System	13
	2.2.1	High Level Model	13
	2.2.2	Example: Robot Modeling	13
	2.2.3	Atomic Propositions and System Executions	14
	2.2.4	Kripke Structure	15
2.3	Specif	ication of the Property to check	16
	2.3.1	Linear-Time Temporal Logic (LTL)	16
	2.3.2	LTL and Kripke structure	18
	2.3.3	Stutter-Invariant LTL Formulas	18
2.4	From	LTL to Büchi Automata	19
	2.4.1	The "Traditional Büchi" Automata (BA)	19
	2.4.2	Generalized Büchi Automata (GBA)	20
	2.4.3	Transition-based Generalized Büchi Automata (TGBA)	23
	2.4.4	Determinization of Büchi Automata	25
	2.4.5	Bisimulation Reduction of Büchi Automata	25
	2.4.6	Converting LTL formula into Büchi Automaton	28
2.5	Explic	it LTL Model checking using Büchi Automata	30
	2.5.1	Synchronous Product	31
	2.5.2	On-the-fly Emptiness check algorithms	33
	2.5.3	Complexity	36
26	Conch	usion	36

3	Eva	uation	of the Testing Automata Approach	39			
	3.1	Introdu	action	39			
	3.2	Stutter	-invariant Languages	40			
	3.3	Testing	g Automata (TA): a natural way to monitor the stuttering	40			
	3.4	3.4 TA Construction					
		3.4.1	From BA to 0-TA: Construction of an intermediate 0-TA from a Büchi				
			Automaton BA	44			
		3.4.2	From $\emptyset$ -TA to TA: Elimination of useless stuttering transitions ( $\emptyset$ ) and				
			introducing livelock-acceptance	45			
		3.4.3	TA Optimizations (that are not yet implemented)	46			
	3.5	Explic	it Model checking using TA	47			
		3.5.1	Synchronous Product of a TA with a Kripke structure	47			
		3.5.2	A two-pass emptiness check algorithm	48			
	3.6	Experi	mental Comparison of TA versus TGBA and BA	52			
		3.6.1	Implementation on top of Spot	52			
		3.6.2	Benchmark Inputs	53			
		3.6.3	Results	54			
		3.6.4	Discussion (TA two-pass emptiness check problem)	61			
	3.7	Conclu	1sion	64			
4	Imp	roving (	the Testing Automata Approach	69			
	4.1	Introdu	uction	69			
	4.2	Improving the Emptiness check by avoiding the second pass in particular cases .					
	4.3	Conve	rting a TA into a Single-pass Testing Automaton (STA)	70			
		4.3.1	Single-pass Testing Automata (STA)	71			
		4.3.2	Construction of an STA from a TA	71			
		4.3.3	Correctness of the one-pass emptiness check using STA	72			
		4.3.4	STA optimization	73			
	4.4	Experi	mental evaluation of the TA improved emptiness check and of STA	75			
		4.4.1	Implementation	75			
		4.4.2	Results	76			
		4.4.3	Discussion	84			
	4.5	Conclu	ision	85			
5	Trai	Transition-based Generalized Testing Automata (TGTA): A Single-pass and Gener-					
	alize	d New	Automata	87			
	5.1	Introdu	action	87			
	5.2	Transit	tion-based Generalized Testing Automata (TGTA)	88			
		5.2.1	Ø-TGTA	89			
		5.2.2	TGTA	90			

	5.3	TGTA Construction	91
		5.3.1 From TGBA to 0-TGTA: Construction of an intermediate 0-TGTA from a	
		TGBA	91
		5.3.2 From Ø-TGTA to TGTA: Elimination of useless stuttering-transitions (Ø)	
		without introducing livelock-acceptance	93
	5.4	Explicit Model checking using TGTA	96
		5.4.1 Synchronous Product of a TGTA with a Kripke Structure	96
		5.4.2 Emptiness check (the same as TGBA)	97
	5.5	Experimental evaluation of TGTA	98
		5.5.1 Implementation	98
		5.5.2 Results	99
		5.5.3 Discussion	107
		5.5.4 Experimental Results once the TGBA is improved by simulation-reduction	108
	5.6	Conclusion	108
п	тти	ing TGTA to improve Symbolic/Hybrid Model Checking 1	111
11		ing 101A to improve Symbolic/Hybrid Woder Checking	
6	Sym	bolic LTL Model Checking using TGTA	113
	6.1	Introduction	113
	6.2	Symbolic LTL Model Checking	114
		6.2.1 Symbolic Kripke Structure	114
		6.2.2 Symbolic Büchi Automata TGBA	115
		6.2.3 Symbolic Product of a TGBA with a Kripke structure	115
		6.2.4 Symbolic Emptiness Check algorithm	116
	6.3	TGTA-based Symbolic LTL Model Checking	117
		6.3.1 Symbolic TGTA	117
		6.3.2 Naive Symbolic Product of TGTA with a Kripke structure	117
		6.3.3 Adjusting the Symbolic Transition Relation of the Kripke Structure to TGTA	118
		6.3.4 Exploiting stuttering transitions to Improve Saturation in the TGTA Ap-	
		proach	119
	6.4	Experimental evaluation	120
		6.4.1 Implementation	122
		6.4.2 Using ETF to build the transition relation of a changeset-based symbolic	
		Kripke structure	122
		6.4.3 Benchmark	124
		6.4.4 Results	124
	6.5	Conclusion	130
7	Hyb	rid LTL Model Checking using TGTA	133
	7.1	Introduction	133
	7.2	Preliminaries	134
		7.2.1 TGBA labeled with propositional formulas	135

		7.2.2	TGTA labeled with propositional formulas	136
	7.3	Symbo	blic Observation Graph (SOG)	140
		7.3.1	SOG	141
		7.3.2	SOG for TGTA (SOG-TGTA)	142
	7.4	Symbo	blic Observation Product (SOP)	143
		7.4.1	SOP	143
		7.4.2	SOP Using TGTA (SOP-TGTA)	144
	7.5	Self-L	oop Aggregation Product (SLAP)	146
		7.5.1	SLAP	146
		7.5.2	SLAP Using TGTA (SLAP-TGTA)	147
	7.6	Experi	mental Comparison of Hybrid Approaches using TGBA vs. TGTA	148
		7.6.1	Implementation	149
		7.6.2	Results	149
		7.6.3	SOG versus SOG-TGTA	150
		7.6.4	SOP versus SOP-TGTA	151
		7.6.5	SLAP versus SLAP-TGTA	152
	7.7	Conclu	ision	153
8	Con	clusion	and Perspectives	155
	8.1	Contex	۰ «t	155
	8.2	Existin	ng Work	155
	8.3	Contri	butions	156
	8.4	Perspe	ctives	158
		8.4.1	Improving TGTA-based approaches	158
		8.4.2	Finding sub-classes of LTL formulas for which TGTA is always efficient	159
		8.4.3	Combining TGTA with Partial Order Reductions	161
Bi	bliogı	raphy		163
A	Exp	eriment	tal Comparison of Explicit approaches using TGBA, TA and TGTA, with	th

-		-
	TGBA improved using simulation-reduction	173

## CHAPTER 1 Introduction

Contents			
1.1	Conte	xt	1
	1.1.1	Model Checking and Automata-theoretic Approach	1
	1.1.2	The State-space Explosion Problem	2
	1.1.3	Explicit versus Symbolic model checking	3
	1.1.4	LTL and Stutter-invariance	3
1.2	Scope	of the Thesis	4
	1.2.1	Existing Work	4
	1.2.2	Contributions	5

## 1.1 Context

Software and hardware systems have become ubiquitous in our everyday life. These systems replace humans for critical tasks that involve high costs and sometimes human lives. This is the case in many areas such as medical devices, telesurgery, nuclear power plants, aircrafts industry, transportation, ...

The serious consequences caused by the failure of such systems make crucial the use of rigorous methods for system validation.

Methods based on testing and simulation have long been used for the validation of systems. However, these techniques allow to explore only a part of the possible system behaviors. The formal verification techniques are exhaustive, they guarantee that a property is satisfied by all possible system executions.

#### 1.1.1 Model Checking and Automata-theoretic Approach

Formal verification [24] provides mathematical-based methods to ensure the correctness of a system with respect to specified behavioral properties (for example, a typical property to check in concurrent systems is the absence of deadlocks). One of the widely-used formal verification methods is *model checking* [25, 54, 8].

Taking as input a high level model describing all possible executions of the system and the property to be checked expressed as a temporal logic formula, a model checker answers if the model satisfies or not the formula. When the property is not satisfied, the model checker returns a *counterexample*, i.e., an execution of the model invalidating the property. This counterexample



Figure 1.1: Automata-theoretic approach to model checking.

is useful to find errors in complex systems. This is an advantage of model checking compared to the other formal methods, such as theorem proving, which can disapprove a property but without providing such a counterexample. Another advantage is the fact that the model checking procedure is completely automatic and easy to use. This automatic procedure is based on the exploration of the system's *state-space*, i.e., a structure that describes all reachable states of the system and all transitions that the system can make between those states. The main disadvantage of model checking is discussed in Section 1.1.2.

The *automata-theoretic approach* [89, 90] to model checking represents this state-space and the property to check using variants of  $\omega$ -automata [42], i.e., an extension of the classical finite automata to recognize words having infinite length (called  $\omega$ -words).

The *automata-theoretic approach* splits the verification process into four operations as shown in figure 1.1:

- 1. Computation of the state-space for the model M. This state-space can be seen as an  $\omega$ -automaton  $A_M$  whose language,  $\mathscr{L}(A_M)$ , represents all possible infinite executions of M.
- 2. Translation of the temporal property  $\varphi$  into an  $\omega$ -automaton  $A_{\neg\varphi}$  whose language,  $\mathscr{L}(A_{\neg\varphi})$ , is the set of all infinite executions that would invalidate  $\varphi$ .
- 3. Synchronization of these automata. This constructs a product automaton  $A_M \otimes A_{\neg \varphi}$  whose language,  $\mathscr{L}(A_M) \cap \mathscr{L}(A_{\neg \varphi})$ , is the set of executions of *M* invalidating  $\varphi$ .
- 4. Emptiness check of this product. This operation tells whether A<sub>M</sub> ⊗ A<sub>¬φ</sub> accepts an infinite word, and can return such a word (a counterexample) if it does. The model *M* verifies φ *iff* L(A<sub>M</sub> ⊗ A<sub>¬φ</sub>) = Ø.

#### 1.1.2 The State-space Explosion Problem

The main difficulty of model checking is the state-space explosion problem [87] caused by the large size of the state-space of the model. For instance, modeling a system of concurrent pro-

cesses can generate a state-space which grows exponentially in the number of processes. Thus, the obtained state-space contains a very important number of states even if each process of the system has only few states. Consequently, in the automata-theoretic approach, the synchronous product of the model's state-space with the property automaton is often too large to be emptiness checked in a reasonable run time and memory. In the literature, different approaches have been proposed to improve the performance of model checking in order to push away the barrier of the State-space Explosion Problem.

The objective of this thesis is to join the fight against this problem, in the context of the verification of stutter-invariant LTL properties. In this work, we consider only finite-state systems (i.e. having a finite number of different configurations). The state-space of these systems have a finite number of states, but possibly infinite number of infinite executions.

#### 1.1.3 Explicit versus Symbolic model checking

There are two main variants of the automata-theoretic approach: explicit and symbolic.

- In the explicit approach [e.g., 26, 48], the state-space of the model and the product are explicitly constructed by enumerating their states. "On-the-fly" emptiness check algorithms avoid the construction of the entire product and state-space by building them lazily during exploration. These on-the-fly algorithms are more efficient because they stop as soon as they find a counterexample and therefore possibly before building the entire product, thereby reducing the amount of memory and time used by the emptiness check in the case of violated properties.
- The symbolic approach [17] tries to overcome the state-space explosion obstacle by representing the state-space implicitly by means of Binary Decision Diagrams (BDDs). The intersection and union of sets of states are translated into conjunction (∧) and disjunction (∨) of Boolean functions efficiently performed using BDDs (or any other variants of Decision Diagrams).

#### 1.1.4 LTL and Stutter-invariance

In order to describe behavioral properties of complex systems, in addition to propositional logic, we need to introduce temporal operators. In this work, we focus on the model checking of stutter-invariant [39] LTL properties.

*Linear Temporal Logic* (LTL), introduced by Pnueli in 1977, is a propositional temporal logic widely used to express temporal properties.

The LTL syntax combines standard logical operators  $(\land, \lor, \neg, \rightarrow, \leftrightarrow)$  and temporal operators to specify that some property *p* happens *next time* "X *p*", *eventually* "F *p*", *always* "G *p*",...(see Section 2.3.1 for a full definition of LTL syntax and semantics).

Typical examples of LTL properties have the form " $G \neg p$ " (*p* never happens), "G F p" (*p* happens infinitely often), and "F G p" (at some point, *p* will hold forever).

LTL properties are constructed over the set *AP* of *atomic propositions*, which represent the properties of individual states. Each state of the system is labeled by a *valuation*  $\ell$  that assigns a truth value to each atomic proposition of *AP* (formally, a valuation is a function  $\ell : AP \mapsto \{\bot, \top\}$ ).

A valuation can also be viewed as a set  $\ell \in 2^{AP}$  interpreted as the set of atomic propositions that are true.

An LTL property  $\varphi$  is interpreted over each *execution* of a system, where an execution maps each time instant to a set of atomic propositions that hold at that instant. Formally, an execution is an infinite sequence of valuations  $\ell_0 \ell_1 \dots \ell_i \dots \in (2^{AP})^{\omega}$ .

The language of  $\varphi$ , denoted  $\mathscr{L}(\varphi) \subseteq (2^{AP})^{\omega}$ , is the set of all sequences of  $(2^{AP})^{\omega}$  satisfying  $\varphi$ . Among LTL properties, we want to distinguish those that are stutter-invariant.

An LTL property  $\varphi$  is *stutter-invariant iff* any sequence  $\ell_0 \ell_1 \ell_3 \dots \in \mathscr{L}(\varphi)$  remains in  $\mathscr{L}(\varphi)$  after repeating any valuation  $\ell_i$  or omitting duplicate valuations. Formally,  $\varphi$  is stutter-invariant iff

$$\ell_0 \ell_1 \ell_2 \dots \in \mathscr{L}(\mathbf{\varphi}) \iff \ell_0^{n_0} \ell_1^{n_1} \ell_2^{n_2} \dots \in \mathscr{L}(\mathbf{\varphi}) \text{ for any } n_0 > 0, n_1 > 0 \dots$$
  
where  $\ell_i^{n_i}$  is the concatenation of  $n_i$  copies of  $\ell_i$ 

Intuitively, stuttering-transitions correspond to transitions that do not change the valuation of atomic propositions between two successive states. Adding or removing stuttering in the system does not change the truth value of stutter-invariant properties. Thus, sequences that differ only in the amount of stuttering can be considered equivalent when checking stutter-invariant properties.

It is well known that any LTLX formula (i.e., an LTL formula that does not use the X operator) describes a stutter-invariant property. Conversely any stutter-invariant property can be expressed as an LTLX formula [70].

According to many research [60, 39], the restriction to stutter-invariant properties is not a serious disadvantage. In addition, there are many tools that enable specific optimizations for the verification of stutter-invariant properties, such as the partial order reduction [86, 69, 50] in Spin tool [55].

#### **1.2** Scope of the Thesis

The general objective of this work is to tackle the state-space explosion problem by improving the performance of the model checking of stutter-invariant LTL properties.

To achieve this goal, we propose some contributions to essentially reduce the size of the product automaton and the computation time/memory used in the emptiness check of this product. Firstly, we start by looking for a form of automata that is suitable for the representation of stutterinvariant properties. As solutions, we propose new types of  $\omega$ -automata that represent all the stuttering-transitions using only self-loops. Then, using these new automata, we propose some contributions to improve the performance of model checking in three contexts: explicit, symbolic and hybrid approaches (where hybrid means combining explicit and symbolic approaches).

#### 1.2.1 Existing Work

Different kinds of  $\omega$ -automata have been used in the automata-theoretic approach to explicit model checking. In the most common case, the property to be checked expressed as an LTL formula is converted into a *Büchi automaton* (BA) [16] with state-based accepting.

In Spot [64], our model checking library, we prefer to represent properties using *general-ized* (i.e., multiple) Büchi acceptance conditions *on transitions* rather than on states, the obtained

automata being called *Transition-based Generalized Büchi Automata* (TGBA). We use TGBA because they allow to have a smaller [49, 36] property automaton than BA.

Unfortunately, having a smaller property automaton  $A_{\neg\varphi}$  does not always imply a smaller product with the model ( $A_M \otimes A_{\neg\varphi}$ ). Thus, instead of targeting smaller property automata, some people have attempted to build automata that are *more deterministic* [78].

Hansen et al. [52, 46] introduced an alternative type of  $\omega$ -automata called *Testing Automata* (TA). TA are like BA, but instead of running them synchronously to the given transition system model, TA only observe changes on the atomic propositions. These automata are less expressive than Büchi automata since they are tailored to represent *stutter-invariant* properties. They are often larger than their equivalent BA, but according to Geldenhuys and Hansen [46], thanks to their high degree of determinism [52], the TA allow to obtain a smaller product and thus improve the performance of model checking. As a back-side, TA have two different modes of acceptance (Büchi-accepting or livelock-accepting), and consequently their emptiness check requires two passes [46], mitigating the benefits of a having a smaller product.

#### **1.2.2** Contributions



Figure 1.2: Yellow colored boxes are original contributions: the new types of  $\omega$ -automata constructed in this work, with the references to the chapters that describe how to use these new automata to improve the model checking of stutter-invariant LTL properties, in three contexts: explicit, symbolic and hybrid approaches.

An overview of the existing automata and the new automata proposed in this work is shown in Figure 1.2.

On the left of this figure, the white boxes show existing  $\omega$ -automata used to represent the property automaton  $A_{\neg\varphi}$ : the traditional *Büchi Automata* (BA) and its generalized variant *Transitionbased Generalized Büchi Automata* (TGBA), and (at the center of the figure) the Testing Automata (TA), the alternative kind of automata that represent only stutter-invariant properties. TA is represented by a box with a double line edge because it requires a two-pass emptiness check algorithm (see Section 3.5.2).

The arrows between the different boxes are labeled by the successive steps used to build the different automata.

In the right part of the figure, the yellow boxes show the new kinds of  $\omega$ -automata constructed in this thesis. Firstly, we propose an improvement of TA called *Single-pass Testing Automata* (STA) (Chapter 4). Secondly, we propose a more efficient new automata (Chapter 5), called *Transition-based Generalized Testing Automata* (TGTA) that combine the advantages of both STA and TGBA, and without the disadvantages of STA.

Then, using these TGTA, several improvements are proposed in three different contexts: explicit model checking (Chapter 5), symbolic model checking based on the saturation algorithm (Chapter 6) and three hybrid techniques (Chapter 7).

These contributions (yellow boxes) are briefly presented in the following:

**Evaluation and improvement of the Testing Automata Approach.** In Chapter 3, We experimentally evaluate the TA approach in order to extend the study of Geldenhuys and Hansen [46]. We show that while TA are statistically more efficient than BA and TGBA when the property to be verified is violated (i.e., a counterexample is found). This is not the case when the property is satisfied since the entire product has to be visited twice to check for each accepting mode of a TA (Büchi-accepting or livelock-accepting). Then, in Chapter 4 we improve the TA approach in two ways. First, we introduce some optimizations on the emptiness check algorithm in order to detect the cases where the second pass is not required. Second, we propose Single-pass Testing Automata (STA), a transformation of TA into a normal form requiring only a single pass during the emptiness check of the product. Although STA are more constrained than TA, we can automatically translate the latter into the former, by adding an artificial livelock-accepting state in STA. We have implemented these improvements in Spot library. We are thus able to compare them with the "traditional" algorithms we used on Testing Automata (TA) and Transition-based Generalized Büchi Automata (TGBA). These experiments show that STA compete well on our examples.

**Transition-based Generalized Testing Automata:** A Single-pass and Generalized New Automata. In Chapter 5, we propose (our main) new type of  $\omega$ -automata for stutter-invariant LTL properties, called *Transition-based Generalized Testing Automata* (TGTA).

TGTA mixes features from both TA and TGBA, without the disadvantage of TA, which is the second pass of the emptiness check, and without adding an artificial state as in STA.

- From TA, TGTA reuses the labeling of transitions with changesets, and the elimination of the useless stuttering-transitions, but without introducing a second mode of acceptance (i.e, livelock-acceptance).
- From TGBA, TGTA inherit the use of transition-based generalized acceptance conditions and reuse the same single-pass emptiness check algorithm.

In addition to improving the performance, the removal of the second pass in TGTA approach also eases the implementation of the emptiness check algorithm, not only for the explicit approach, but especially for more complex implementations (such as symbolic or hybrid approaches). More generally, this simplification eases the combination of TGTA with other classical optimizations used in model checking (such as the *partial order reduction*, or the saturation technique used in the symbolic approach).

Another advantage of TGTA compared to TA is that a TA is built from a BA while a TGTA

is built from a TGBA. Therefore, TGTA can take advantage from the fact that TGBA are more concise [49, 36] than BA.

Compared to TGBA, TGTA represents all stuttering-transitions with only self-loops on all states (thanks to the elimination of the useless stuttering-transitions during the TGTA construction).

Implementation and experimentation of TGTA approach show that in most cases, it reduces the size of the synchronous product and is statistically more efficient than TA and TGBA (BA) approaches, for the explicit model checking of stutter-invariant properties,

Using TGTA to improve a Saturation-based Symbolic Model Checking. In symbolic model checking, the product is symbolically encoded by means of decision diagrams [17] and computed as a least fixpoint on its symbolic transition relation. The performance of this fixpoint computation can be improved using the saturation technique [20, 85]. In order to improve the symbolic model checking of stutter-invariant LTL properties, we investigate in Chapter 6 the use of the combination of TGTA with saturation technique. We first show how a TGTA can be symbolically encoded, Then, we show that the saturation algorithm greatly benefits from the presence of stuttering self-loops on all states of TGTA, and we propose a symbolic encoding of stuttering transitions in the product that improve the saturation-based symbolic approach using TGTA.

Implementation and experimentation of this approach confirm that it outperforms the saturation-based symbolic approach using Büchi Automata TGBA, the performance of the saturation algorithm are significantly more enhanced by TGTA than by TGBA.

This improvement was possible only because TGTA represents the stuttering-transitions specifically in a way that helps the saturation technique.

**Three hybrid approaches using TGTA** The *hybrid approaches* [79, 38] combine ideas from both explicit and symbolic approaches in order to benefit from the advantages of the both worlds, i.e., encoding the set of states in a concise way using decision diagrams as in the symbolic approach, and the emptiness check performed on-the-fly as in the explicit approach.

In this work, we focus on three hybrid techniques proposed in [37]: the *Symbolic Observation Graph* (SOG), the *Symbolic Observation Product* (SOP) and the *Self-Loop Aggregation Product* (SLAP).

As for symbolic model checking, Testing Automata have never been used before for hybrid model checking.

The three hybrid approaches SOG, SOP and SLAP are based on TGBA. In Chapter 7, we define and implement variations of these three approaches using TGTA instead of TGBA. Then, each original approach (SOG, SOP and SLAP) is experimentally compared against its TGTA variant (respectively SOG-TGTA, SOP-TGTA and SLAP-TGTA). The results show that the performance of these new variants depend on the type of the formula to be checked: verified or violated.

Part I

Preliminaries

## CHAPTER 2

# The Traditional Approaches to LTL Model Checking

#### Contents

2.1	Introduction 1				
2.2	2 Modeling the System				
	2.2.1	High Level Model	13		
	2.2.2	Example: Robot Modeling	13		
	2.2.3	Atomic Propositions and System Executions	14		
	2.2.4	Kripke Structure	15		
2.3	Specif	fication of the Property to check	16		
	2.3.1	Linear-Time Temporal Logic (LTL)	16		
	2.3.2	LTL and Kripke structure	18		
	2.3.3	Stutter-Invariant LTL Formulas	18		
2.4	From	LTL to Büchi Automata	19		
	2.4.1	The "Traditional Büchi" Automata (BA)	19		
	2.4.2	Generalized Büchi Automata (GBA)	20		
	2.4.3	Transition-based Generalized Büchi Automata (TGBA)	23		
	2.4.4	Determinization of Büchi Automata	25		
	2.4.5	Bisimulation Reduction of Büchi Automata	25		
	2.4.6	Converting LTL formula into Büchi Automaton	28		
2.5	Explic	tit LTL Model checking using Büchi Automata	30		
	2.5.1	Synchronous Product	31		
	2.5.2	On-the-fly Emptiness check algorithms	33		
	2.5.3	Complexity	36		
2.6	Concl	usion	36		

## 2.1 Introduction

The Model checking of a behavioral property on a finite-state system is an automatic procedure that requires many phases. The first step is to formally represent the system and the property to be checked. The formalization of the system produces a model M that formally describes all



Figure 2.1: Automata-theoretic Approach to model checking.

the possible behaviors of the system. The property to be checked is formally described using a specification language such as branching-time (CTL) or Linear-time Temporal Logic (LTL). In this work we concentrate on LTL. The next step is to run a model checking algorithm that takes as inputs the model *M* and the LTL formula  $\varphi$ . This algorithm exhaustively checks that all the model *M* behaviors satisfy  $\varphi$ .

There are two major approaches for LTL model checking: explicit and symbolic. The *explicit* [e.g., 26, 48] approach converts M and  $\varphi$  into explicit graphs before running the model checking algorithm. The symbolic approach [17] encodes M and  $\varphi$  implicitly by means of Binary Decision Diagrams (BDDs). This symbolic approach will be presented in Chapter 6, we only consider the explicit approach in this chapter.

The automata-theoretic approach to explicit model checking relies on  $\omega$ -automata, i.e., an extension of finite automata to infinite words.

Figure 2.1 summarizes the successive steps of the automata-theoretic approach. It starts by converting the negation of  $\varphi$  into an  $\omega$ -automaton  $A_{\neg\varphi}$ , then composing that automaton with the state-space of a model *M* given as a Kripke structure  $\mathcal{K}_M$  (a variant of  $\omega$ -automaton), and finally checking the language emptiness of the resulting product automaton  $A_{\neg\varphi} \otimes \mathcal{K}_M$  [89].

As for any model checking process, the automata-theoretic approach suffers from the well known state explosion problem [87]. In practice, it is the product automaton that can be very large, its size can reach  $(|A_{\neg \varphi}| \times |\mathcal{K}_M|)$  states, which can make it impossible to be handled using the resources of modern computers.

The  $\omega$ -automaton representing  $A_{\neg\varphi}$  is usually a Büchi automaton (BA) or a generalization using multiple acceptance sets, such as Generalized Büchi Automata (GBA) or Transition-based Generalized Büchi Automata (TGBA).

This chapter details the successive phases of the model checking procedure, including the formalization and the different algorithms used in these phases. We also present the different variants of Büchi automata and their use in the automata-theoretic approach to model checking.

## 2.2 Modeling the System

A model checker tool performs verification on a model of the system rather than the system itself. A model is a high-level representation that reproduces the relevant part of the behaviors of the original system, while eliminating irrelevant details that are difficult to reproduce. The advantages are that the model includes only the relevant elements of the entire larger system, and the model is easier to build and to redesign once possible errors are reported by the model checker.

#### 2.2.1 High Level Model

The model of a system can be described using several high-level formalisms such as Petri nets [31], Promela programs [55], communicating finite-state machines [15], or finite transition systems [1]. All these formalisms can be used to generate the state-space of the model, but the difficulty for some formalisms is to prove that the state-space is finite. For example, a state-space generated from a Petri net can be infinite, because it corresponds to the graph of reachable markings [31]. To ensure that this state-space is finite, it is necessary to check that the Petri net places are bounded.

Whatever the used formalism, a finite state-space can be viewed as a form of finite transition system, called *Kripke Structure* (its formal definition is given later).

#### 2.2.2 Example: Robot Modeling

Let us consider a robot specification [9] as an illustrative example of System modeling.



Figure 2.2: A model for robot behavior

In the finite transition system of Figure 2.2, the robot can execute only three actions: *look*, *compute*, and *move*. It begins by taking a snapshot of its environment, this action is represented by the transition *look* executed from the initial state labeled by "*Ready to look*". Then, from the state "*Ready to compute*" the robot executes the transition *compute* to determine its future location and finally it moves to this computed location by executing the transition *move* from the state "*Ready to move*". However, if the robot does not find a new location in the current snapshot of its environment, it returns to the initial state to take a new snapshot, hoping that its environment has changed.

#### 2.2.3 Atomic Propositions and System Executions

The *atomic propositions* allow to describe the different states of the system. The set of atomic propositions is the set of smallest properties defined for individual states of the system (where smallest means that the truth value of each atomic proposition does not depend on the truth value of the other atomic propositions).

For instance, in order to describe the different states of the robot of Figure 2.2, let us define the three following atomic propositions:

- l = The robot is ready to look
- c = The robot is ready to compute
- m = The robot is ready to move



Figure 2.3: A Kripke structure for the robot model presented in Figure 2.2.

Each system state is labeled by a *valuation*, i.e., an assignment of truth value to each atomic proposition of *AP*.

**Definition 1** (Valuation). Let AP a finite set of atomic propositions, a valuation  $\ell$  over AP is represented by a function  $\ell : AP \mapsto \{\bot, \top\}$ .

We denote by  $\Sigma = 2^{AP}$  the set of all valuations over AP, where a valuation  $\ell \in \Sigma$  is interpreted either as the set of atomic propositions that are true, or as a Boolean conjunction where:

$$\ell$$
 is identified to  $(\bigwedge_{p\in\ell}p)\wedge (\bigwedge_{p\in(AP\setminus\ell)}ar{p})$ 

For instance if  $AP = \{a, b\}$ , then  $\Sigma = 2^{AP} = \{\{a, b\}, \{a\}, \{b\}, \emptyset\}$  or equivalently  $\Sigma = \{ab, a\bar{b}, \bar{a}b, \bar{a}\bar{b}\}$  such that:

$\{a,b\}$	$\leftrightarrow$	$a \wedge b$	( <i>ab</i> )
$\{a\}$	$\leftrightarrow$	$a \wedge \bar{b}$	$(a\overline{b})$
$\{b\}$	$\leftrightarrow$	$\bar{a} \wedge b$	$(\bar{a}b)$
Ø	$\leftrightarrow$	$\bar{a}\wedge \bar{b}$	$(\bar{a}\bar{b})$

For the robot model, the set of valuations is  $\Sigma = 2^{\{l,c,m\}}$ . The initial state is labeled by the valuation  $\{l\}$  also noted  $(l \wedge \bar{c} \wedge \bar{m})$  or  $(l\bar{c}\bar{m})$ , the valuations of the other states are shown in Figure 2.3.

**Definition 2** (Sequence of Valuations). Let  $n \in \mathbb{N} \cup \{\omega\}$  where  $\omega$  is the lowest transfinite ordinal number *defined by Cantor. A* sequence  $\sigma$  *of n valuations of*  $\Sigma$  *is a function*  $\sigma : [[0, n[] \mapsto \Sigma$  *mapping each index from* [[0, n[] *to a valuation of*  $\Sigma$ .

In the following, a sequence  $\sigma$  is often represented by the concatenation of its valuations:  $\sigma = \sigma(0) \cdot \sigma(1) \cdot \sigma(2) \cdots$ .

 $\Sigma^n$  = the set of sequences of length *n* and  $\Sigma^* = \bigcup_{n \in \mathbb{N}} \Sigma^n$  the set of finite sequences of valuations from  $\Sigma$ .

 $\Sigma^{\omega}$  denotes the set of infinite sequences of valuations from  $\Sigma$ .

**Definition 3** (Execution). An execution of a system maps to each time instant, a set of atomic propositions that hold at that instant. This execution is represented by an infinite sequence of valuations  $\sigma = \sigma(0) \cdot \sigma(1) \cdot \sigma(2) \cdots \in \Sigma^{\omega}$ .

The set of all executions of a system is a subset of  $\Sigma^{\omega}$  and can be viewed as a language of infinite words over the alphabet  $\Sigma$ .

An example of an execution of the robot system is the sequence of valuations  $\{l\} \cdot \{c\} \cdot \{l\} \cdot \{c\} \cdot \{l\} \cdot \{c\} \cdots$  (i.e., the robot execution alternating the valuations  $\{l\}$  and  $\{c\}$ ), we can notice that there is a risk that the robot never reaches the state labeled by the valuation  $\{m\}$  and therefore never moves.

In the next section, we present a variant of transition system usually used in model checking process to describe the system executions.

#### 2.2.4 Kripke Structure

The state-space of a system can be represented by a directed graph, called Kripke structure, where vertices represent the states of the system and edges are the transitions between these states. In addition, each vertex is labeled by a valuation that represents the set of atomic propositions that are true in the corresponding state.

**Definition 4** (Kripke Structure). A Kripke structure over the set of atomic propositions AP is a tuple  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$ , where:

- *S* is a finite set of states,
- $S_0 \subseteq S$  is the set of initial states,
- $\mathcal{R} \subseteq S \times S$  is the transition relation,
- *l*: S → Σ is a labeling function that maps each state s to a valuation that represents the set of atomic propositions that are true in s.

An infinite path or a run of a Kripke structure  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$  is an infinite sequence of states  $r = s_0 \cdot s_1 \cdot s_2 \cdots$  such that  $s_0 \in S_0$  and  $\forall i \in \mathbb{N}, (s_i, s_{i+1}) \in \mathcal{R}$ .

*We denote*  $\operatorname{Run}(\mathcal{K})$  *the set of infinite paths of*  $\mathcal{K}$ *.* 

Given an infinite path  $r = s_0 \cdot s_1 \cdot s_2 \cdots$  in  $\operatorname{Run}(\mathcal{K})$ , the infinite sequence of valuations  $\sigma = l(s_0) \cdot l(s_1) \cdot l(s_2) \cdots \in \Sigma^{\omega}$  corresponds to an execution of the system represented by  $\mathcal{K}$ . Thus,  $\sigma$  is also called *execution* of  $\mathcal{K}$ .

**Definition 5.** The language of a Kripke structure  $\mathcal{K}$  is the set  $\mathscr{L}(\mathcal{K}) \subseteq \Sigma^{\omega}$  of all executions of  $\mathcal{K}$ .

$$\mathscr{L}(\mathscr{K}) = \{ \mathbf{\sigma} = l(s_0) \cdot l(s_1) \cdot l(s_2) \cdots \in \Sigma^{\omega} \mid s_0 \cdot s_1 \cdot s_2 \cdots \in \operatorname{Run}(\mathscr{K}) \}$$

Note that  $\mathscr{L}(\mathscr{K}) \subseteq \Sigma^{\omega}$  can be viewed as a language of infinite words over the alphabet  $\Sigma = 2^{AP}$ .

Figure 2.3.2 shows the Kripke structure representing the state-space of the robot system (described in section 2.2.2).

## 2.3 Specification of the Property to check

In Model checking, it is necessary to have a precise formal expression of properties to check. Logic can express these properties in a mathematical, unambiguous, and concise way. This enables the automation of the verification.

#### 2.3.1 Linear-Time Temporal Logic (LTL)

In order to express temporal properties, *Linear Temporal Logic* (LTL) [62] adds the notion of causality to the traditional propositional logic. LTL introduces new operators, called temporal operators, interpreted over linear executions. In a linear execution, every time instant has only one immediate successor, unlike the other widely-used temporal logic, called CTL (*Computational Tree Logic*) [23], for which a time instant can have many immediate successors. This two logics are not really comparable [55]. Each logic allows to express properties that the other can not express. For instance, CTL allows to express the reset [55] properties and LTL cannot. However, CTL does not allow to express the invariance LTL properties (of the form FGp). In this work, we only focus on LTL model checking.

#### 2.3.1.1 LTL Formula

An LTL formula is composed of:

- A finite set of atomic propositions  $AP = \{p_1, p_2, \ldots\},\$
- the Boolean operators  $\land$ ,  $\lor$ ,  $\neg$ ,  $\rightarrow$ ,  $\leftrightarrow$ ,
- the unary temporal operators X (next), G (Globally) and F (Future),
- the binary temporal operator U (Until).

**Definition 6.** *Given a set of atomic propositions AP. The set of LTL formulas over AP is inductively defined as follows:* 

- For every atomic proposition  $p \in AP$ , p is an LTL formula,
- *if*  $\phi$  *is an LTL formula, then*  $\neg \phi$ , X $\phi$ , G $\phi$  *and* F $\phi$  *are also LTL formulas,*
- *if*  $\varphi_1$  *and*  $\varphi_2$  *are two LTL formulas, then*  $(\varphi_1 \land \varphi_2)$ *,*  $(\varphi_1 \lor \varphi_2)$ *,*  $(\varphi_1 \to \varphi_2)$ *,*  $(\varphi_1 \leftrightarrow \varphi_2)$  *and*  $(\varphi_1 \cup \varphi_2)$  *are also LTL formulas.*

#### 2.3.1.2 LTL Formula Semantics

An LTL formula over *AP* is interpreted over an execution  $\sigma \in \Sigma^{\omega}$  ( $\Sigma = 2^{AP}$ ).

In the following,  $\sigma(n)$  denotes the  $(n+1)^{th}$  valuation of the execution  $\sigma = \sigma(0)\sigma(1)\sigma(2)...$ , and the sequence  $\sigma^i = \sigma(i)\sigma(i+1)...$  denotes the suffix of  $\sigma$  starting at position *i* (i.e.,  $\forall n \in \mathbb{N}, \sigma^i(n) = \sigma(i+n)$ ).

**Definition 7.** The satisfaction of an LTL formula  $\varphi$  by an execution  $\sigma \in \Sigma^{\omega}$ , denoted  $\sigma \models \varphi$ , is defined inductively as follows:

 $\begin{aligned} \sigma &\models p & \text{iff } p \in \sigma(0) \\ \sigma &\models \neg \phi & \text{iff } \neg(\sigma \models \phi) \\ \sigma &\models \phi_1 \land \phi_2 & \text{iff } \sigma \models \phi_1 \text{ and } \sigma \models \phi_2 \\ \sigma &\models \phi_1 \lor \phi_2 & \text{iff } \sigma \models \phi_1 \text{ or } \sigma \models \phi_2 \\ \sigma &\models X\phi & \text{iff } \sigma^1 \models \phi \ (\phi \text{ is true in the next time step}) \\ \sigma &\models G\phi & \text{iff } \forall i \ge 0, \sigma^i \models \phi \ (\phi \text{ is true in every time step}) \\ \sigma &\models F\phi & \text{iff } \exists i \ge 0 \text{ such that } \sigma^i \models \phi \ (\phi \text{ is true now or at some future time step}) \\ \sigma &\models \phi_1 \cup \phi_2 & \text{iff } \exists i \ge 0 \text{ such that } \sigma^i \models \phi_2 \text{ and } \forall j \in [[0, i-1]], \sigma^j \models \phi_1 \ (\phi_2 \text{ is true now or } \phi_1 \text{ is true now and } \phi_1 \text{ remains true until } \phi_2 \text{ holds}) \end{aligned}$ 

**Definition 8.** Given an LTL formula  $\varphi$  over AP. The language of  $\varphi$ , denoted  $\mathscr{L}(\varphi)$ , is the set of all executions  $\sigma \in \Sigma^{\omega}$  satisfying  $\varphi$ .

$$\mathscr{L}(\boldsymbol{\varphi}) = \{ \boldsymbol{\sigma} \in \boldsymbol{\Sigma}^{\boldsymbol{\omega}} \mid \boldsymbol{\sigma} \models \boldsymbol{\varphi} \}$$

#### 2.3.1.3 LTL Operators Equivalences

In Definition 6, some operators are redundant (i.e., can be expressed using other operators). The following equivalences allow to ignore these operators in the LTL translation algorithm presented in Section 2.4.6.

$$T = True = p \lor \bar{p}$$

$$\bot = False = p \land \bar{p}$$

$$F\phi = T \cup \phi$$

$$G\phi = \neg F \neg \phi = \neg (T \cup \neg \phi) = \bot R\phi$$

#### 2.3.1.4 Size of an LTL formula

**Definition 9.** Given an LTL formula  $\varphi$  over AP, the size (or the length) of  $\varphi$ , denoted  $|\varphi|$ , is the total number of symbols of  $\varphi$ , i.e., the number of atomic propositions, constants, and operators (logical and temporal) occurring in  $\varphi$ .

Formally, the size of an LTL formula is defined inductively as follows:

$$\begin{split} |\bot| &= 1 \\ |\top| &= 1 \\ |p| &= 1, \forall p \in AP \\ |\neg \phi| &= 1 + |\phi| \\ \\ |\phi_1 \land \phi_2| &= |\phi_1 \lor \phi_2| = 1 + |\phi_1| + |\phi_2| \\ \\ X \phi| &= |G \phi| = |F \phi| = 1 + |\phi| \\ \\ |\phi_1 \cup \phi_2| &= 1 + |\phi_1| + |\phi_2| \end{split}$$

#### 2.3.2 LTL and Kripke structure

The link between an LTL formula and a Kripke structure is obvious, since we can interpret the LTL formula over the executions of the Kripke structure .

**Definition 10.** We say that a Kripke structure  $\mathcal{K}$  over  $\Sigma = 2^{AP}$ , satisfies an LTL formula  $\varphi$ , denoted  $\mathcal{K} \models \varphi$ , if all executions in  $\mathscr{L}(\mathcal{K})$  satisfy  $\varphi$ .

$$\mathcal{K} \models \varphi \iff \mathscr{L}(\mathcal{K}) \subseteq \mathscr{L}(\varphi)$$

In the example of the robot model, the Kripke structure of Figure does not satisfy the LTL formula  $\varphi = Fm$ : a counterexample (i.e., an execution of the Kripke structure that does not satisfy  $\varphi$ ) is  $\{l\} \cdot \{c\} \cdot \{l\} \cdot \{c\} \cdot \{l\} \cdot \{c\} \cdots$  (the robot alternating  $\{l\}$  and  $\{c\}$  and never reaches the state labeled by the valuation  $\{m\}$ ; therefore it never moves).

#### 2.3.3 Stutter-Invariant LTL Formulas

**Definition 11.** An LTL formula  $\varphi$  is stutter-invariant [39] iff any sequence  $\sigma(0)\sigma(1)... \in \mathscr{L}(\varphi)$ remains in  $\mathscr{L}(\varphi)$  after repeating any valuation  $\sigma(i)$  or omitting duplicate valuations. Formally,  $\varphi$ is stutter-invariant iff

$$\sigma(0)\sigma(1)\ldots\in\mathscr{L}(\varphi)\iff\underbrace{\sigma(0)\ldots\sigma(0)}_{n_0\times\sigma(0)}\underbrace{\sigma(1)\ldots\sigma(1)}_{n_1\times\sigma(1)}\ldots\in\mathscr{L}(\varphi), \text{ for any } n_0>0, n_1>0\ldots$$

Intuitively, a property is stutter-invariant if it is insensitive to stuttering transitions i.e., the transitions that do not change the values of atomic propositions between two states of the system.

For example, the LTL formula  $\varphi = a \cup b$  is stutter-invariant because valuations can be repeated or omitted in any sequence that satisfies  $\varphi$  (e.g.,  $a\bar{b};\bar{a}b;\bar{a}b;\bar{a}b;\bar{a}\bar{b}\ldots$ ), and the obtained sequences remain satisfying  $\varphi$  (e.g.,  $a\bar{b};\bar{a}b;\bar{a}\bar{b}\ldots$ ). Conversely, the LTL formula  $\varphi = a \wedge Xb$  is not stutter-invariant because it is satisfied by the sequence  $(a\bar{b};\bar{a}b;\ldots)$  but unsatisfied by the sequence  $(a\bar{b};\bar{a}b;\ldots)$ .

**Theorem 1.** Any  $LTL \setminus X$  formula (i.e., an LTL formula that does not use the X (next-time) operator) describes a stutter-invariant property. Conversely, any stutter-invariant property can be expressed as an  $LTL \setminus X$  formula [70].

### 2.4 From LTL to Büchi Automata

The automata-theoretic approach is based on the transformation of the LTL formula to check into an automaton that accepts the same executions (called infinite words in the context of automata theory).

The following sections present three variants of Büchi automata [16] that can be used to express properties in the automata-theoretic approach to model checking: the standard Büchi Automata (BA), Generalized Büchi Automata (GBA) and Transitions-based Generalized Büchi Automata (TGBA). The main difference between these automata is the way they accept an infinite word of  $(2^{AP})^{\omega}$ .

#### 2.4.1 The "Traditional Büchi" Automata (BA)

The Büchi Automata were introduced by J.R Büchi [16] in 1962. They are  $\omega$ -automata [42] with labels on transitions and acceptance conditions on states. While classical finite automata recognize words having finite length,  $\omega$ -automata (and in particular Büchi Automata) recognize words of infinite length, called  $\omega$ -words. Although they accept infinite words, Büchi Automata have a finite number of states.

In the following, we will use the abbreviation BA for the standard variant of Büchi Automata.

**Definition 12** (BA). A Büchi Automaton (BA) over the alphabet  $\Sigma = 2^{AP}$  is a tuple  $\mathcal{B} = \langle Q, I, \delta, \mathcal{F} \rangle$  where:

- *Q* is a finite set of states,
- $I \subseteq Q$  is a a finite set of initial states,
- $\mathcal{F} \subseteq Q$  is a finite set of accepting states ( $\mathcal{F}$  is called the accepting set),
- $\delta \subseteq Q \times \Sigma \times Q$  is the transition relation where each transition is labeled by a letter  $\ell$  of  $\Sigma$ , i.e., each element  $(q, \ell, q') \in \delta$  represents a transition from state q to state q' labeled by a valuation  $\ell \in 2^{AP}$ .

**Definition 13** (A BA run). A run of  $\mathcal{B}$  over an infinite word  $\sigma = \ell_0 \ell_1 \ell_2 \dots \in \Sigma^{\omega}$  is an infinite sequence of transitions  $r = (q_0, \ell_0, q_1)(q_1, \ell_1, q_2)(q_2, \ell_2, q_3) \dots \in \delta^{\omega}$  such that  $q_0 \in I$  (i.e., the infinite word is recognized by the run). Such a run is said to be accepting if  $\forall i \in \mathbb{N}, \exists j \ge i, q_j \in \mathcal{F}$  (i.e., at least one accepting state is visited infinitely often).

The infinite word  $\sigma$  is accepted by  $\mathcal{B}$  if there exists an accepting run of  $\mathcal{B}$  over  $\sigma$ . The language accepted by  $\mathcal{B}$  is the set  $\mathscr{L}(\mathcal{B}) \subseteq \Sigma^{\omega}$  of the infinite words it accepts.

Figure 2.4 shows two examples of BA:

- (a) Figure 2.4a is a BA recognizing the runs where *a* is true infinitely often and *b* is true infinitely often, i.e., recognizing the LTL formula  $(GFa \wedge GFb)$ ,
- (b) Figure 2.4b is a BA recognizing the LTL formula  $(a \cup Gb)$ .

In these BA, the initial states are indicated by an arrow without source state " $\rightarrow$ ". The Boolean conjunctions labeling each transition are valuations over  $AP = \{a, b\}$ . The accepting states are indicated by a double circle. The LTL formulas labeling each state represent the property accepted



Figure 2.4: Two examples of BA, with accepting states shown as double circles.

starting from this state of the automaton: they are shown for the reader's convenience but not used for model checking.

As an illustration of the BA definition, the infinite word  $ab; a\bar{b}; \bar{a}b; ab; \bar{a}b; ab; ...$  is accepted by the BA of Figure 2.4b that recognizes  $a \cup Gb$ . A run over such infinite word must start in the initial state labeled by the formula  $(a \cup Gb)$  and remains in this state for the first two valuations  $ab; a\bar{b}$ , then it changes the value of a, so it has to take the transition labeled by the valuation  $\bar{a}b$  to move to the second state labeled by the formula (Gb). Finally, to be accepted, it must stay on this accepting state by executing infinitely the transitions labeled by  $\{\bar{a}b, ab\}$ . The obtained accepting run is:  $(a \cup Gb) \xrightarrow{ab} (a \cup Gb) \xrightarrow{\bar{a}b} (a \cup Gb) \xrightarrow{\bar{a}b} (Gb) \xrightarrow{\bar{a}b} (Gb) \xrightarrow{\bar{a}b} (Gb) \cdots (Gb) \cdots$ 

#### 2.4.2 Generalized Büchi Automata (GBA)

In the generalized variant of Büchi Automata (GBA) [48], there are multiple acceptance conditions, in other words,  $\mathcal{F}$  is a set of accepting sets of states and a run is accepted *iff* it visits infinitely often each accepting set in  $\mathcal{F}$ .

**Definition 14** (GBA). A Generalized Büchi Automata (GBA) over the alphabet  $\Sigma = 2^{AP}$  is a tuple  $\mathcal{G} = \langle Q, I, \delta, \mathcal{F} \rangle$  where:

- *Q* is a finite set of states,
- $I \subseteq Q$  is a set of initial states,
- $\mathcal{F} \subseteq 2^{\mathcal{Q}}$  is a set of sets of accepting states (we call accepting set each set  $F_i$  of  $\mathcal{F} = \{F_1, F_2, \ldots, F_k \ldots\}$ ),
- $\delta \subseteq Q \times \Sigma \times Q$  is the transition relation where each transition is labeled by a letter  $\ell \in \Sigma$ (i.e., each element  $(q, \ell, q') \in \delta$  represents a transition from state q to state q' labeled by a valuation  $\ell \in 2^{AP}$ ).

**Definition 15.** A run of  $\mathcal{G}$  over an infinite word  $\sigma = \ell_0 \ell_1 \ell_2 \ldots \in \Sigma^{\omega}$  is an infinite sequence of transitions  $r = (q_0, \ell_0, q_1)(q_1, \ell_1, q_2)(q_2, \ell_2, q_3) \ldots \in \delta^{\omega}$  such that  $q_0 \in I$  (i.e., the infinite word is recognized by the run).

Such a run is said to be accepting if  $\forall F \in \mathcal{F}, \forall i \ge 0, \exists j \ge i, q_j \in F$  (i.e., at least one accepting state of each accepting set  $F \in \mathcal{F}$  is visited infinitely often).

The infinite word  $\sigma$  is accepted by G iff there exists an accepting run of G over  $\sigma$ . The language accepted by G is the set  $\mathscr{L}(G) \subseteq \Sigma^{\omega}$  of infinite words it accepts.



(a) A GBA for  $\varphi = GFa \wedge GFb$ 

Figure 2.5: Two examples of GBA: (a) A GBA for  $\varphi = GFa \wedge GFb$  with two accepting sets of states indicated by  $\bullet$  and  $\circ$ . (b) A GBA for  $a \cup Gb$  with a single accepting set indicated by  $\bullet$ .

GBA are as expressive as BA: any GBA can be converted into an equivalent BA that recognizes the same language (section 2.4.2.1), and vice-versa, a BA  $\langle Q, I, \delta, \mathcal{F} \rangle$  can be viewed as a GBA  $\langle Q, I, \delta, \{\mathcal{F}\} \rangle$  having a single set of accepting states  $\{\mathcal{F}\}$ .

Figure 2.5 shows two examples of GBA that are equivalent to the two BA of Figure 2.4 such that:

- (a) The GBA of Figure 2.5a recognizes (GFa∧GFb). An accepting run in this GBA has to visit the two accepting states indicated by and infinitely often. Therefore, it must explore infinitely often, the transition labeled by {ab, ab} (a is true) and the transition {āb, ab} (b is true).
- (b) The GBA of Figure 2.5b is the same as the BA of Figure 2.4b, with the accepting state is indicated by the single acceptance condition  $\{\bullet\}$ . More generally, a BA  $\langle Q, I, \delta, \mathcal{F} \rangle$  can be viewed as a GBA with a single set of accepting states  $\langle Q, I, \delta, \mathcal{F} \rangle$ .

#### 2.4.2.1 From GBA to BA (degeneralization)

The conversion of a GBA into a BA is called *degeneralization* [4, section 4.3.4]. A GBA having *s* states and *k* accepting sets of states can be *degeneralized* into an equivalent BA with  $(k \times s)$  states.

Given a GBA  $\mathcal{G} = \langle Q, I, \delta, \{\mathcal{F}_1, \mathcal{F}_2, \dots, \mathcal{F}_k\} \rangle$ , the BA  $\mathcal{B} = \langle Q', Q'^0, \delta', \mathcal{F}' \rangle$  constructed as follows accepts the same language as  $\mathcal{G}$ .

- $Q' = Q \times \{1, ..., k\}$  (Q' is the states of Q marked by an integer in  $\{1, ..., k\}$ )
- $Q'^0 = Q \times \{1\}$
- $\mathcal{F}' = \mathcal{F}_1 \times \{1\}$
- $\forall (q,i) \in Q', ((q,i),\ell,(q',i')) \in \delta' \text{ if } (q,\ell,q') \in \delta \text{ and } \begin{cases} i'=i & \text{if } q \notin \mathcal{F}_i \\ i'=(i \mod k)+1 & \text{if } q \in \mathcal{F}_i \end{cases}$  (In the

 $i^{th}$  copy, the states of  $\mathcal{F}_i$  are connected to its successors in the  $(i+1)^{th}$  copy (the  $(k+1)^{th}$  copy is the first copy)

Figure 2.6 illustrates the successive steps to degeneralize the GBA G of  $GFa \wedge GFb$ :

The GBA G is presented in Figure 2.6a with F = {F<sub>1</sub>, F<sub>2</sub>} = {{q<sub>0</sub>}, {q<sub>1</sub>}} respectively indicated by ● and ○.



(a) GBA  $\mathcal{G}$  for  $\varphi = (\mathsf{GF}a \land \mathsf{GF}b)$ , with  $\mathcal{F} = \{\{q_0\}, \{q_1\}\}$  such that  $\mathcal{F}_1 = \{q_0\}$  and  $\mathcal{F}_2 = \{q_1\}$  respectively indicated by  $\bullet$  and  $\bigcirc$ .



(b) Duplicate the GBA (two copies)



(d) Computing the transitions of the second copy: i.e., states  $(q_0, 2)$  and  $(q_1, 2)$ 



(c) Computing the transitions of the first copy: i.e., states  $(q_0, 1)$  and  $(q_1, 1)$ 



(e) The obtained BA  $\mathcal{B}$  for  $\varphi = (\mathsf{GF}a \land \mathsf{GF}b)$ , with the accepting state is a double circle

Figure 2.6: The successive steps to degeneralize a GBA into BA.

- The degeneralization begin by duplicate k times the GBA, with  $k = |\mathcal{F}| = 2$ . The obtained BA of Figure 2.6b is composed of two copies of the GBA, the initial states are only in the first copy.
- In Figure 2.6c, we add the transitions of the first copy of states, the state (q<sub>0</sub>, 1) is connected to its successors in the second copy because q<sub>0</sub> ∈ F<sub>1</sub>.
- Similarly, in Figure 2.6d, because q<sub>1</sub> ∈ F<sub>2</sub>, the state (q<sub>1</sub>,2) of the second copy is connected to its successors in the first copy. In the two copies, we keep unchanged the transitions of the two other states (q<sub>1</sub>,1) and (q<sub>0</sub>,2).
- Finally, the obtained BA  $\mathcal{B}$  is shown in Figure 2.6e, it could be reduced by merging the two bisimilar states  $q'_1$  and  $q'_3$ . This optimization is called bisimulation reduction and will be presented in section 2.4.5.
- Each accepting run of *CB* is also accepted by *G*:

An accepting run  $r_1$  of the obtained BA  $\mathcal{B}$  has to visit infinitely often at least one accepting state that is in the first copy. To achieve this,  $r_1$  must visit all the intermediate copies that are only connected through accepting set of the GBA (each  $\mathcal{F}_i$  connects the  $i^{th}$  copy to the  $(i+1)^{th}$  copy). Therefore,  $r_1$  has to visit each accepting set  $\mathcal{F}_i$  infinitely often, thus  $r_1$  is also an accepting run for the original GBA  $\mathcal{G}$ .

• Conversely, each accepting run of CG is also accepted by  $\mathcal{B}$ :

We can consider, without loss of generality, that an accepting run  $r_2$  of the GBA  $\mathcal{G}$  visits infinitely often all the accepting set in the order  $\mathcal{F}_1, \mathcal{F}_2, \ldots$ . Therefore,  $r_2$  visits infinitely often at least one state of  $\mathcal{F}' = \mathcal{F}_1 \times \{1\} = \{(q_0, 1)\} = \{q'_0\}$  with  $q'_0$  is an accepting state of  $\mathcal{B}$ , thus  $r_2$  is also an accepting run for the BA  $\mathcal{B}$ .

#### 2.4.3 Transition-based Generalized Büchi Automata (TGBA)

A Transition-based Generalized Büchi Automaton (TGBA) [49] is a variant of a Büchi automaton that has multiple acceptance conditions on transitions.

**Definition 16** (TGBA). A TGBA over the alphabet  $\Sigma = 2^{AP}$  is a tuple  $\mathcal{G}' = \langle Q, I, \delta, \mathcal{F} \rangle$  where:

- *Q* is a finite set of states,
- $I \subseteq Q$  is a set of initial states,
- $\mathcal{F}$  is a finite set of acceptance conditions,
- $\delta \subseteq Q \times \Sigma \times 2^{\mathcal{F}} \times Q$  is the transition relation, where each element  $(q, \ell, F, q') \in \delta$  represents a transition from state q to state q' labeled by a valuation  $\ell \in 2^{AP}$ , and a set of acceptance conditions  $F \in 2^{\mathcal{F}}$ .

**Definition 17.** A run of G' over an infinite word  $\sigma = \ell_0 \ell_1 \ell_2 \dots \in \Sigma^{\omega}$  is an infinite sequence of transitions  $r = (q_0, \ell_0, F_0, q_1)(q_1, \ell_1, F_1, q_2)(q_2, \ell_2, F_2, q_3) \dots \in \delta^{\omega}$  such that  $q_0 \in I$  (i.e., the infinite



Figure 2.7: (a) A TGBA with acceptance conditions  $\mathcal{F} = \{\bullet, \circ\}$  recognizing the LTL property  $\phi = \mathsf{GF}a \land \mathsf{GF}b$ . (b) A TGBA with  $\mathcal{F} = \{\bullet\}$  recognizing the LTL property  $a \mathsf{UG}b$ .

word is recognized by the run). Such a run is said to be accepting if  $\forall f \in \mathcal{F}, \forall i \in \mathbb{N}, \exists j \ge i, f \in F_j$ (i.e., each acceptance condition is visited infinitely often).

The infinite word  $\sigma$  is accepted by G' iff there exists an accepting run of G' over  $\sigma$ . The language of G' is the set  $\mathscr{L}(G') \subseteq \Sigma^{\omega}$  of infinite words it accepts.

TGBA and BA have the same expressive power: any TGBA can be converted into a languageequivalent BA and vice-versa [27, 49]. The process of converting a TGBA into a BA is also called *degeneralization* and is similar to the "GBA to BA transformation" presented in section 2.4.2.1.

Figure 2.7 shows the same properties as Figure 2.4 and Figure 2.5, but expressed as TGBA. A run in these examples is accepted if it visits infinitely often all acceptance conditions (represented by colored dots  $\circ$  and  $\bullet$  on transitions).

- (a) The TGBA of Figure 2.7a recognizes the runs that contain infinitely often *a* and infinitely often *b*. An accepting run in this TGBA has to visit infinitely often the two acceptance conditions indicated by and ●. Therefore, it must explore infinitely often the transitions where *a* is true (i.e., transitions labeled by *ab* or *ab*) and infinitely often the transitions where *b* is true (i.e., transitions labeled by *ab* or *ab*). As an illustration of the degeneralization process, the BA from Figure 2.4a was built by degeneralizing the TGBA from Figure 2.7a. In the worst case, a TGBA with *s* states and *n* acceptance conditions will be degeneralized into a BA with *s* × (*n*+1) states. The worst case of the degeneralization occurred here, since the TGBA with 1 state and *n* acceptance conditions was degeneralized into a BA with *n*+1 states. It is known that no BA with less than *n*+1 states can accept the property \(\lambda\_{i=1}^n GF\_{p\_i}\) so this BA is optimal [22] in terms of number of states.
- (b) The property *a*UG*b* is easier to express: the TGBA of Figure 2.7b is the same as the BA of Figure. 2.4b, with the acceptance condition moved on transitions. More generally, a BA can be seen as a TGBA, by simply marking as accepting the transitions leaving the accepting states, without adding states nor transitions. Algorithms that input TGBA can therefore be easily adapted to process BA or GBA.

#### 2.4.4 Determinization of Büchi Automata

**Definition 18.** A Büchi automaton is deterministic iff each accepted infinite word of  $\Sigma^{\omega}$  is recognized by an unique accepting run of  $\delta^{\omega}$ .

Formally, a Büchi automaton  $\langle Q, I, \delta, \mathcal{F} \rangle$  is deterministic *iff* |I| = 1 and  $\forall (q, \ell) \in Q \times \Sigma$ ,  $|\{q' \in Q \mid (q, \ell, q') \in \delta\}| \leq 1$  (i.e., the outgoing transitions of each state are labeled with different valuations).

As an illustration, the Büchi automaton of Figure 2.4a is a deterministic BA. However, the Büchi automaton of Figure 2.4b is a non-deterministic BA.

The definition 18 of deterministic Büchi Automata is valid for the three variants of Büchi Automata: BA, GBA and TGBA.

The determinization operation consists in the transformation of a non-deterministic Büchi automaton into a deterministic one. However, there are non-deterministic Büchi automata that cannot be determinized. For instance, the non-deterministic Büchi automata of Figure 2.8 that recognize the LTL property  $\varphi = FGa$ , cannot be transformed into an equivalent deterministic Büchi automaton.

The classical powerset construction [72] used to determinize finite automata does not work for any Büchi automata (it works only for a restricted class of Büchi automata [30]).

The powerset construction uses subsets of states of the non-deterministic automaton, as states of the deterministic automaton. The example of Figure 2.9 shows that this construction does not work for the Büchi automata of  $\varphi = FGa$ . The constructed automaton of Figure 2.9b accepts the infinite word  $a; \bar{a}; a; \bar{a}; ...$  (alternating a and  $\bar{a}$  infinitely) and therefore does not recognize the same language as the original automaton of Figure 2.9a.

Thus, the determinization of Büchi automata requires the use of other types of  $\omega$ -automata, such as Muller or Rabin automata [76].

The history of determinizing Büchi automata began in 1963 with the first attempt of determinization due to Muller [65] that appears to be faulty. In 1966, McNaughton [63] proved that any non-deterministic Büchi automaton can be converted into a deterministic Muller automaton. This conversion involves a doubly exponential blow-up in the size of the original Büchi automaton  $(2^{2^{O(n)}})$ .

Proposed in 1988, the Safra's construction [76] transforms a non-deterministic Büchi automaton with n states into an equivalent deterministic Muller or Rabin automaton with  $2^{O(n \log(n))}$  states (optimal for Rabin automata).

To conclude this section, it is important to note that although there are LTL formulas that are not convertible into Büchi automaton and the cost to obtain a deterministic automaton remains very high, we will see in the following that even the use of automata that are not "completely" deterministic but having a high degree of determinism can improve the performance of model checking.

#### 2.4.5 Bisimulation Reduction of Büchi Automata

We say that two states are bisimilar if the automaton can accept the same runs starting from either of these states (this implies that two bisimilar states recognize the same language). Büchi Au-



(a) BA/GBA for  $\varphi = FGa$ 

(b) TGBA for  $\varphi = FGa$ 

Figure 2.8: Examples of non-deterministic Büchi Automata: BA/GBA and TGBA for the LTL property  $\phi = FGa$ . Acceptance states/transitions are indicated by  $\bullet$ .



Figure 2.9: Büchi automata determinization: The classical powerset construction does not work for  $\phi = FGa$  (A deterministic Büchi automaton for FGa does not exist).

tomata can be simplified by merging bisimilar states. This bisimulation reduction can be achieved using a partition refinement algorithm [e.g., 14, 40, 13, 91, 88].

The main idea of this algorithm is to split the set of states of the automaton into equivalence classes according to the equivalence relation of bisimilarity, then the states of each equivalence class are merged into a single state.

One way to compute the equivalence classes is to use the notion of signature [13]. A signature sig(q) can be viewed as a "fingerprint" of each state q that encodes the outgoing transitions of q. Two states q and q' having different signatures are not bisimilar and therefore belong to different classes C and C'.

- For BA ⟨Q, I,δ, F⟩, the signature of a state q with respect to a partition Π, is the set of pairs (valuation, destination class C ∈ Π) of each outgoing transition from q: i.e., sig<sub>Π</sub>(q) = {(ℓ,C) | ∃q' ∈ C ∈ Π, (q,ℓ,q') ∈ δ}.
- For TGBA, the signature also includes the acceptance conditions of transitions:  $sig_{\Pi}(q) = \{(\ell, F, C) \mid \exists q' \in C \in \Pi, (q, \ell, F, q') \in \delta\}.$
- In the case of a GBA with  $\mathcal{F} = \{F_1, F_2, \dots, F_n\}$ , the signature must includes the acceptance conditions on states:  $sig_{\Pi}(q) = \{(\ell, F, C) \mid \exists q' \in C \in \Pi, (q, \ell, F, q') \in \delta\}$ , where  $F = \{i \mid q \in F_i\}$ .

The sets representing the signatures can be encoded symbolically [91, 14], BDDs (Binary Decision Diagrams) are used to implement the signatures of the bisimulation reduction of the different automata used in this work.

The basic idea of the bisimulation algorithm is to build a partition of the equivalence classes by refining an initial partition  $\Pi_0$ . This can be implemented by the following iterative procedure :

- Initialization:
  - For a standard Büchi Automaton BA, we set up with two equivalence classes  $\mathcal{F}$  and  $Q \setminus \mathcal{F}$ :  $\Pi_0 = \{\mathcal{F}, Q \setminus \mathcal{F}\}.$
  - For a TGBA or a GBA,  $\Pi_0 = \{Q\}$ .
- Iterate until a fixpoint is reached (i.e.,  $\Pi_k = \Pi_{k+1}$ ):
  - For each pair of states q and q' belonging to the same class C of the partition  $\Pi_k$ , if  $sig_{\Pi_k}(q) \neq sig_{\Pi_k}(q')$  then q and q' are put into two different classes C' and C'' of the new partition  $\Pi_{k+1}$

When this procedure stops, the states of each class have the same signature in the obtained partition. The termination of this algorithm is guaranteed because the number of states is finite and therefore the number of partition too. This naive implementation has a quadratic complexity (due to the comparison of all the pairs of states signatures, for each iteration). The optimization [88] reduces this complexity to  $O(m\log(n))$ , where *n* is the number of states and *m* is the number of transitions.



(a) A BA  $\mathcal{B}$  before bisimulation reduction

(b)  $\mathcal{B}$  after bisimulation reduction

Figure 2.10: An example of Bisimulation Reduction of BA

As an illustration, Figure 2.10a and Figure 2.10b show an example of BA respectively before and after applying the bisimulation procedure:

The successive iterations to compute the equivalence classes from the automaton of Figure 2.10a are:

- $\Pi_0 = \{\{q_0, q_1, q_2, q_3, q_4\}, \{q_5\}\}$ , because  $q_5$  is the unique accepting state.
- $\Pi_1 = \{\{q_0, q_1, q_3\}, \{q_2, q_4\}, \{q_5\}\}, \text{ because in } \Pi_0, \text{ we have } (\bar{a}b, \{q_5\}) \in sig_{\Pi_0}(q_2) = sig_{\Pi_0}(q_4) \text{ and } (\bar{a}b, \{q_5\}) \notin sig_{\Pi_0}(q_0) = sig_{\Pi_0}(q_1) = sig_{\Pi_0}(q_2).$

- $\Pi_2 = \{\{q_0\}, \{q_1, q_3\}, \{q_2, q_4\}, \{q_5\}\}$ , because in  $\Pi_1$ , we have  $(\bar{a}b, \{q_2, q_4\}) \in sig_{\Pi_1}(q_1) = sig_{\Pi_1}(q_3)$  and  $(\bar{a}b, \{q_2, q_4\}) \notin sig_{\Pi_1}(q_0)$ .
- $\Pi_3 = \Pi_2$  the procedure stops.

The automaton obtained after this partitions refinement is shown in Figure 2.10b.

Simulation [40, 41, 44, 81, 11] reduction is a generalization and an improvement of the bisimulation reduction. Simulation relation between two states q and q' is based on the inclusion of the sets of infinite runs starting from q and q', instead of equality between these sets of infinite runs as in bisimulation relation between q and q'. Recently, [3] proposed a simulation reduction that improves the determinism of the resulting automaton.

#### 2.4.6 Converting LTL formula into Büchi Automaton

Any LTL formula  $\varphi$  can be converted into a Büchi Automaton whose language is the set of executions that satisfy  $\varphi$ . This conversion is central to the process of model checking, thus a substantial number of research has been conducted in this area. Several algorithms have been proposed to translate an LTL formula into a BA, GBA or TGBA.

The first translation algorithm was proposed by Wolper [93, 92], it allows to convert an LTL formula  $\varphi$  into an automaton whose size is always equal to  $2^{O(|\varphi|)}$  states (exponential in the size of  $\varphi$ ). In [92], the algorithm proposed is dedicated to LTL and produces a GBA. Despite its exponential complexity, this algorithm is easier to understand than others. In the following, we present a concise variant of this algorithm proposed in [66].

Using the equivalences between the LTL operators (Section 2.3.1.3), we can assume that the LTL formula  $\varphi$  to translate, is only composed of the operators  $\neg$ ,  $\lor$ , X and U. There are several optimizations based on rewriting rules[81, 40, 11] that simplify the LTL formula before translating it into an automaton.

We note  $A_{\varphi}$  the GBA constructed by the following translation. The computation of the states of  $A_{\varphi}$  is based on the set  $cl(\varphi)$ , the *closure* of  $\varphi$  containing all the subformulas of  $\varphi$  and  $X(\varphi_1 \cup \varphi_2)$  for each subformula  $(\varphi_1 \cup \varphi_2)$  of  $\varphi$ , it is in addition closed by the operator  $\neg$ .

**Definition 19** (Closure). Let  $sub(\varphi)$  be the set of all subformulas of  $\varphi$ . The closure of  $\varphi$  is the smallest set  $cl(\varphi)$  satisfying:

- $sub(\mathbf{\varphi}) \subseteq cl(\mathbf{\varphi})$ ,
- {X( $\phi_1 \cup \phi_2$ ) | ( $\phi_1 \cup \phi_2$ )  $\in sub(\phi)$ }  $\subseteq cl(\phi)$ ,
- $\forall \phi \in cl(\phi), (\neg \phi) \in cl(\phi)$  (*With reduction of redundancies like*  $\phi = \neg \neg \phi$  *and*  $\neg X \phi = X \neg \phi$ ).

For example,  $cl(\mathsf{F}p) = cl(\top \mathsf{U}p) = \{p, \bar{p}, (\top \mathsf{U}p), \mathsf{X}(\top \mathsf{U}p), \neg(\top \mathsf{U}p), \mathsf{X}\neg(\top \mathsf{U}p)\}$ . It is easy to deduce that  $|cl(\varphi)| \in O(|\varphi|)$ .

Each state of  $A_{\varphi}$  is labeled by an *atom*, where an atom is a subset of  $cl(\varphi)$  that satisfies the following three consistency rules:

**Definition 20** (Atom). An atom  $\alpha$  of  $\varphi$  is a subset of  $cl(\varphi)$ , satisfying the following rules:

• Logical consistency rules ( $\alpha$  is maximal and does not contain logical contradictions):

-  $\forall \phi \in cl(\phi), \phi \in \alpha \Leftrightarrow \neg \phi \notin \alpha$ , this rule also implies that  $\alpha$  is maximal, i.e.,  $\forall \phi \in cl(\phi)$ , either  $\phi \in \alpha$  or  $\neg \phi \in \alpha$ ,
$- \forall (\phi_1 \lor \phi_2) \in cl(\phi), (\phi_1 \lor \phi_2) \in \alpha \Leftrightarrow (\phi_1 \in \alpha) \text{ or } (\phi_2 \in \alpha),$ 

• Temporal consistency rule:

 $\forall (\phi_1 \cup \phi_2) \in cl(\phi), (\phi_1 \cup \phi_2) \in \alpha \Leftrightarrow (\phi_2 \in \alpha) \text{ or } ((\phi_1 \in \alpha) \text{ and } X(\phi_1 \cup \phi_2) \in \alpha)).$ *This temporal consistency rule is deduced from the expansion law stating that:*  $\phi_1 \cup \phi_2 = \phi_2 \lor (\phi_1 \land X(\phi_1 \cup \phi_2)).$ 

Let  $Atoms(\varphi)$  denotes the set of atoms of  $\varphi$ , we have  $Atoms(\varphi) \subseteq 2^{cl(\varphi)}$  and the number of states of  $A_{\varphi}$  is  $|A_{\varphi}| = |Atoms(\varphi)| \leq |2^{cl(\varphi)}| \in O(2^{|\varphi|})$ .



Figure 2.11: A GBA  $A_{\varphi}$  constructed form ( $\varphi = F p$ ) by the algorithm of Definition 21.

As an illustration, we use the above rules to compute the atoms for the translation of the LTL formula  $\varphi = F p = (\top U p)$ , we obtain the set Atom(F p) composed of the four atoms labeling the states of the GBA  $A_{(Fp)}$  shown in Figure 2.11. This set of atoms is obtained by the following reasoning:

In order to satisfy the temporal consistency rule, the atoms that contain the sub-formula  $(\top \cup p)$  must also contain p or  $X(\top \cup p)$ . Thus, we obtain three atoms containing  $(\top \cup p)$ :

- $\{(\top U p), p, X(\top U p)\} \in Atom(F p)$
- $\{(\top \cup p), \bar{p}, \mathsf{X}(\top \cup p)\} \in Atom(\mathsf{F}p)$
- $\{(\top \cup p), p, X \neg (\top \cup p)\} \in Atom(\mathsf{F} p)$

On the other hand, the atoms that do not contain  $(\top \cup p)$  must contain  $\neg(\top \cup p)$  because they have to be maximal (the first rule), and they can only contain  $\bar{p}$  and  $\neg X(\top \cup p)$  to satisfy the temporal consistency rule. We obtain an unique atom that contain  $\neg(\top \cup p)$ :  $\{\neg(\top \cup p), \bar{p}, X \neg(\top \cup p)\} \in$ *Atom*(F p).

Once the atoms labeling the states computed, we can build the GBA  $A_{\phi}$  according to the following definition:

**Definition 21** (From  $\varphi$  to  $A_{\varphi}$ ). *Given an LTL formula*  $\varphi$  *over AP, A GBA over the alphabet*  $\Sigma = 2^{AP}$  *that accepts the same language as*  $\varphi$ *, is a tuple*  $A_{\varphi} = \langle Q, I, \delta, \mathcal{F} \rangle$  *where:* 

- $Q = Atoms(\varphi)$ ,
- $I = \{ \alpha \in Atoms(\phi) \mid \phi \in \alpha \},\$

- $\delta = \{(\alpha, \alpha \cap AP, \alpha') \in Q \times \Sigma \times Q \mid \forall X \phi \in cl(\phi), (X \phi \in \alpha) \Leftrightarrow (\phi \in \alpha')\},\$
- $\mathcal{F} = \{\mathcal{F}_{\phi_1 \cup \phi_2} \mid (\phi_1 \cup \phi_2) \in cl(\phi)\}, \text{ with } \mathcal{F}_{\phi_1 \cup \phi_2} = \{\alpha \in Q \mid \phi_2 \in \alpha \lor (\phi_1 \cup \phi_2) \notin \alpha\}$

This construction of  $A_{\varphi}$  is illustrated in Figure 2.11 presenting the GBA  $A_{(F_p)}$  obtained from the LTL formula  $\varphi = F_p = (\top \cup p)$ . Intuitively, when the constructed automaton  $A_{\varphi}$  reads a word  $\sigma = \ell_0 \ell_1 \ell_2 \dots$  by exploring a sequence of states  $\alpha_0 \alpha_1 \alpha_2 \dots$ , in each step *i*, the "sub-"word  $w_i = \ell_i \ell_{i+1} \ell_{i+1} \dots$  satisfies all the formulas labeling the state  $\alpha_i$ . In addition,  $w_i$  does not satisfy any other state of  $Q \setminus \alpha_i$  because  $w_i$  does not satisfy any formulas of  $cl(\varphi) \setminus \alpha_i$  by definition of the atom  $\alpha_i$ . The formulas of the form X $\phi$  are verified in the next step by the transition relation  $\delta$ , because  $\delta$ is constructed according to the equivalence  $(X\phi \in \alpha_i) \Leftrightarrow (\phi \in \alpha_{i+1})$  where  $(\alpha_i, \alpha_i \cap AP, \alpha_{i+1}) \in \delta$ (Definition 21).

The automata produced by translation algorithms often contain redundant states, which can be eliminated using the bisimulation/simulation rOneductions presented in Section 2.4.5.

As we mentioned earlier, the algorithm presented above is not optimal, it produces a GBA composed of 4 states (Figure 2.11) for  $\varphi = F p$ , while this LTL formula can be converted into a GBA having only 2 states.

Many other more efficient algorithms are proposed to translate an LTL formula into the different variants of Büchi Automata. One common way to obtain a BA from an LTL formula is to first translate the formula into some Generalized Büchi Automata with multiple acceptance conditions (GBA, TGBA,...) and then to *degeneralize* this automaton to obtain a single acceptance condition. Alternatives include the translation of the property into a *state-based* [48] generalized automaton which can then also be degeneralized, or the translation of the property into an alternating Büchi automaton that is then converted into a BA using the construction proposed by [45].

The degeneralization process can increase the size of the Büchi automaton (see Section 2.4.2.1). In addition, several model checking procedures supports generalized (i.e., multiple) acceptance conditions, making such degeneralization unnecessary and even costly [29]. Moving the acceptance conditions from the states (GBA) to the transitions (TGBA) also reduces the size of the property automaton [27, 49].

Several algorithms exist to translate an LTL formula into a TGBA [49, 27, 28, 84, 45, 2]. The one we use in the experimentations presented in this thesis is based on Couvreur's LTL translation algorithm [27]. The next chapter presents the results of an experimentation that compares different model checking approaches, including those using BA and TGBA.

Most of the works that tried to improve model checking, have focused on translating the LTL formula into the smallest possible automaton. However, [78] claims (without giving evidence) that model checking can be improved by generating more deterministic automata (determinization of Büchi automata was discussed in section 2.4.4). According to [3, 10], it is not yet clear when "more deterministic" automaton should be preferred to a small one.

# 2.5 Explicit LTL Model checking using Büchi Automata

The explicit model checking enumerates the states of the state-spaces of both the model and the LTL formula. These explicit state-spaces and their synchronous product are traditionally repre-

sented using variants of  $\omega$ -automata. In this automata-based approach to model checking of an LTL formula  $\varphi$  on the model *M*, there are two important algorithms (represented in Figure 2.1):

- 1. The translation of the negation of the LTL formula  $\neg \varphi$  into (one variant of) a Büchi Automaton  $A_{\neg\varphi}$  (Translation of LTL properties into automata has already been presented in the previous section).
- 2. The emptiness check of the product  $\mathcal{K}_{M} \otimes A_{\neg \varphi}$  where  $\mathcal{K}_{M}$  is the Kripke structure representing the state-space of the model M. The language of this product automaton  $\mathscr{L}(\mathcal{K}_{M} \otimes A_{\neg \varphi})$  is equal to  $\mathscr{L}(\mathcal{K}_{M}) \cap \mathscr{L}(A_{\neg \varphi})$ , i.e., the set of executions of  $\mathcal{K}_{M}$  invalidating  $\varphi$ .

The goal of the emptiness check algorithm is to determine if the product automaton  $\mathcal{K}_M \otimes A_{\neg \varphi}$  accepts an execution or not.

In other words, it checks if the language of the product automaton is empty or not. If it is empty, then there is no execution of the model *M* that invalidates the property  $\varphi$  and therefore  $M \models \varphi$ . Otherwise, there is an execution of *M* that invalidates  $\varphi$  and this execution is reported as a counterexample.

More formally, the automata-theoretic approach is based on the following equivalences:

$$\begin{aligned} \mathscr{L}(\mathscr{K}_{\mathcal{M}} \otimes A_{\neg \varphi}) &= \emptyset \iff \mathscr{L}(\mathscr{K}_{\mathcal{M}}) \cap \mathscr{L}(A_{\neg \varphi}) = \emptyset \\ \iff \mathscr{L}(\mathscr{K}_{\mathcal{M}}) \cap \overline{\mathscr{L}(A_{\varphi})} = \emptyset \\ \iff \mathscr{L}(\mathscr{K}_{\mathcal{M}}) \subseteq \mathscr{L}(A_{\varphi}) \\ \iff M \models \varphi \end{aligned}$$

The emptiness check algorithms dealing with TGBA also work for GBA and BA, because a GBA or a BA can be seen as a TGBA by pushing the acceptance conditions on the transitions leaving accepting states (with multiple acceptance conditions in the case of GBA). For this reason, this section only focuses on TGBA.

### 2.5.1 Synchronous Product

The product of a TGBA with a Kripke structure is a TGBA whose language is the intersection of both languages.

**Definition 22** (Synchronous Product of a Kripke structure and a Büchi automaton). *For a Kripke structure*  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$  *and a TGBA* 

- $\mathcal{A} = \langle Q, I, \delta, \mathcal{F} \rangle$  the product  $\mathcal{K} \otimes \mathcal{A}$  is the TGBA  $\langle \mathcal{S}_{\otimes}, I_{\otimes}, \delta_{\otimes}, \mathcal{F} \rangle$  where
  - $S_{\otimes} = S \times Q$ ,
  - $I_{\otimes} = \mathcal{S}_0 \times I$ ,
  - $\delta_{\otimes} = \{((s,q),\ell,F,(s',q')) \mid (s,s') \in \mathcal{R}, (q,\ell,F,q') \in \delta, l(s) = \ell\}$

**Property 1.** We have  $\mathscr{L}(\mathcal{K} \otimes \mathcal{A}) = \mathscr{L}(\mathcal{K}) \cap \mathscr{L}(\mathcal{A})$  by construction.



Figure 2.12: Example of a Synchronous Product  $\mathcal{K} \otimes \mathcal{A}$  between a Kripke structure  $\mathcal{K}$  and a TGBA  $\mathcal{A}$  recognizing the LTL formula FG*p*, with acceptance conditions indicated by the black dot  $\bullet$ .

Because the set of all states of the product is defined as  $S_{\otimes} = S \times Q$ , the maximum size (in term of number of reachable states) of the product automaton is equal to  $|S| \times |Q|$ .

Figure 2.12 shows an example of a Synchronous Product between a Kripke structure  $\mathcal{K}$  and a TGBA (BA)  $\mathcal{A}$  recognizing the LTL formula FG p. Each state of  $\mathcal{K}$  is numbered and labeled with the set of atomic propositions (of  $\Sigma = \{p\}$ ) that hold in this state. In the TGBA representing the product  $\mathcal{K} \otimes \mathcal{A}$ , the states are labeled with a pairs of the form "( $\mathcal{K}$  state,  $\mathcal{A}$  state)".

Let us recall that since a BA or a GBA can be seen as a TGBA with accepting conditions moved from accepting states onto their outgoing transitions, the above product construction using TGBA can easily be adapted to a product between a Kripke structure and a BA or a GBA.

### 2.5.2 On-the-fly Emptiness check algorithms

Testing a TGBA for emptiness amounts to the search of an accepting cycle that contains at least one occurrence of each acceptance condition. This can be done in different ways: either with a variation of Tarjan or Dijkstra algorithm [27] or using several Nested Depth-First Searches (NDFS) to save some memory [84]. In NDFS algorithms, a first Depth-First Search (DFS) is performed until it reaches an accepting state *s*, then a second DFS is performed from *s* trying to return to *s*. [29, 46] argued that NDFS algorithms are slower than SCC-based algorithms, so in the following, we will use in our experiments Couvreur's SCC-based emptiness check algorithm [27] because it only performs a single DFS, and its complexity does not depend on the number of acceptance conditions (while NDFS may need to perform many nested DFS in the case of multiple acceptance conditions (GBA or TGBA)). This Couvreur's algorithm will be described in detail below. The comparison of the different emptiness check algorithms has raised many studies [47, 77, 29, 73].

The product automaton that has to be explored during the emptiness check is generally very large, its size can reach the value obtained by multiplying the the sizes of the model and formula automata, which are synchronized to build this product. Therefore, building the entire product must be avoided. "On-the-fly" emptiness check algorithms allow the product automaton to be constructed lazily during its exploration. These on-the-fly algorithms are more efficient because they stop as soon as they find a counterexample and therefore possibly before building the entire product, thereby reducing the amount of memory and time used by the emptiness check.

In this work, we focus on SCC-based "On-the-fly" emptiness checks. Algorithm 1 presented below is an iterative version of the Couvreur's SCC-based algorithm [27].

Algorithm. 1 computes on-the-fly the Maximal Strongly Connected Components (MSCCs) of the TGBA representing the product  $\mathcal{K} \otimes \mathcal{A}$ : it performs a Depth-First Search (DFS) for SCC detection and then merges the SCCs belonging to the same maximal SCC into a single SCC. After each merge, if the union of all acceptance conditions occurring in the merged SCC is equal to  $\mathcal{F}$  (line 16), then an accepting run (i.e., a counterexample) is found (line 16) and the  $\mathcal{L}(\mathcal{K} \otimes \mathcal{A})$  is not empty. *todo* is the DFS stack. It is used by the procedure DFSpush to push the states of the current DFS path and the set of their successors that have not yet been visited. *H* maps each visited state to its rank in the DFS order, and H[s] = 0 indicates that *s* is a dead state (i.e., *s* belongs to a maximal SCC that has been fully explored). Figure 2.14 illustrates a run of this algorithm on a small example.

The SCC stack stores a chain of partial SCCs found during the DFS. For each SCC the attribute

1 **Input**: A product TGBA  $\mathcal{K} \otimes \mathcal{A} = \langle \mathcal{S}_{\otimes}, I_{\otimes}, \delta_{\otimes}, \mathcal{F} \rangle$ **2 Result**:  $\top$  if and only if  $\mathscr{L}(\mathscr{K} \otimes \mathscr{A}) = \emptyset$ **3 Data**: *todo*: stack of  $\langle state \in S, succ \subseteq \delta_{\otimes} \rangle$ *SCC*: stack of  $(root \in \mathbb{N}, la \subseteq \mathcal{F}, acc \subseteq \mathcal{F}, rem \subseteq \mathcal{S})$ *H*: map of  $\mathcal{S} \mapsto \mathbb{N}$  $max \leftarrow 0$ 4 begin foreach  $s^0 \in I_{\otimes}$  do 5 DFSpush( $\emptyset$ ,  $s^0$ ) 6 while ¬todo.empty() do 7 if *todo*.top().*succ* =  $\emptyset$  then 8 DFSpop() 9 else 10 11 pick one  $\langle s, \_, a, d \rangle$  off *todo*.top().*succ* if  $d \notin H$  then 12 DFSpush(*a*, *d*) 13 else if H[d] > 0 then 14 merge(a, H[d]) 15 if  $SCC.top().acc = \mathcal{F}$  then 16 return ot17 return ⊤ 18 19 DFSpush ( $la \subseteq \mathcal{F}$ ,  $s \in \mathcal{S}_{\otimes}$ )  $max \leftarrow max + 1$ 20  $H[s] \leftarrow max$ 21 *SCC*.push( $\langle max, la, \emptyset, \emptyset \rangle$ ) 22 *todo*.push( $\langle s, \{\langle q, l, a, d \rangle \in \delta_{\otimes} \mid q = s\}$ )) 23 24 DFSpop() 25  $\langle s, \_ \rangle \leftarrow todo.pop()$ SCC.top().rem.insert(s) 26 if H[s] = SCC.top().root then 27 foreach  $s \in SCC.top().rem$  do 28  $H[s] \leftarrow 0$ 29 SCC.pop() 30 31 merge ( $la \subseteq \mathcal{F}, t \in \mathbb{N}$ )  $r \leftarrow \emptyset$ 32  $acc \leftarrow la$ 33 while *t* < *SCC*.top().*root* **do** 34  $acc \leftarrow acc \cup SCC.top().acc$ 35  $\cup$  SCC.top().la  $r \leftarrow r \cup SCC.top().rem$ 36 37 SCC.pop()  $SCC.top().acc \leftarrow SCC.top().acc \cup acc$ 38  $SCC.top().rem \leftarrow SCC.top().rem \cup r$ 39

Algorithm 1: Emptiness check algorithm for TGBA.



Figure 2.13: SCC stack: the use of the SCCs fields la and acc.



Figure 2.14: Six intermediate steps in a run of algorithm 1. The states  $s_1, \ldots, s_5$  are labeled by their value in *H*. The stack of roots of SCCs (the *root* stack in the algorithm) and the DFS search stack (induced by the recursive calls to DFSpush()) are displayed on the side. An interpretation of the SCC stack in term of SCCs is given as yellow blobs on the automaton.

(a) Initially the algorithm performs a DFS search by declaring each newly encountered state as a trivial SCC. (b) When the transition from  $s_4$  to  $s_3$  is processed, the algorithm detects that  $H[s_3] \neq 0$  which means the transition creates a cycle and all SCCs between  $s_4$  and  $s_3$  are merged. (c) When the DFS exits the non-accepting  $\{s_3, s_4\}$  SCC, it marks all its states as dead (H[s] = 0). (d) When the DFS attempt to visit a dead state, it ignores it. (e) Visiting the transition from  $s_5$  to  $s_1$  will merge three SCCs into one, but it does not yet appear to be accepting because the white accepting condition ( $\bigcirc$ )has not been seen. (f) Finally visiting the transition from  $s_2$  back to  $s_1$  will add the white acceptance condition to the current SCC, and the algorithm will stop immediately because it has found an SCC labeled by all acceptance conditions.

*root* is the DFS rank (H) of the first state of the SCC, *acc* is the set of all acceptance conditions belonging to the SCC, *la* is the acceptance conditions of the transition between the previous and the current SCC, and *rem* contains the fully explored states of the SCC. Figure 2.13 shows how *acc* and *la* are used in the SCC search stack.

- 1. The algorithm begins by pushing in *SCC* each state visited for the first time (line 4), as a trivial SCC with an empty *acc* set (line 22).
- 2. Then, when the DFS explores a transition t between two states s and d, if d is in the SCC stack (line 11), therefore t closes a cycle passing through s and d in the product automaton. This cycle "strongly connects" all SCCs pushed in the SCC stack between SCC[i] and SCC[n]: the two SCCs that respectively contains the states d and s (SCC[n] is the top of the SCC stack).
- 3. All the SCCs between SCC[i] and SCC[n] are merged (line 15) into SCC[i]. The merge of acceptance conditions is illustrated by Figure 2.13: a "back" transition t is found between SCC[n] and SCC[i], therefore the latest SCCs (from i to n) are merged.
- 4. The acceptance conditions of the merged SCC is equal to the union of SCC[i].acc ∪ SCC[i+1].acc ∪ SCC[i+1].acc ∪ ··· ∪ SCC[n].la ∪ SCC[n].acc ∪ t.la. If this union is equal to F, then the merged SCC is accepting and the algorithm return *false* (line 16): the product is not empty.

Figure 2.14 illustrates the successive steps of a run of algorithm 1.

### 2.5.3 Complexity

In the automata-theoretic approach to model checking of a Kripke structure  $\mathcal{K}_M$  against an LTL formula  $\varphi$ , the upper bound of the time (and space) complexity is in  $O(|\mathcal{K}_M| \times 2^{|\varphi|})$ , because  $\varphi$  can be converted into Büchi automaton  $A_{\neg\varphi}$  whose size and time of construction is in  $O(2^{|\varphi|})$ , and the emptiness check is linear with respect to the size of the product automaton  $\mathcal{K}_M \otimes A_{\neg\varphi}$ . This linear complexity is for instance the complexity of Algorithm 1, which has the same complexity as the algorithm for finding the maximum strong components in a directed graph of Dijkstra [32, 33].

Thus, the upper bound of the model checking complexity is directly deduced from the inequality:  $|\mathcal{K}_{\mathcal{M}} \otimes A_{\neg \varphi}| \leq |\mathcal{K}_{\mathcal{M}}| \times |A_{\neg \varphi}|$ .

# 2.6 Conclusion

Automata-theoretic approach is traditionally used for the explicit LTL model checking. In this approach, a Kripke structure  $\mathcal{K}_M$  is used to represent the state-space of the model M, and the property to be checked is expressed as an LTL formula  $\varphi$ , then its negation is converted into a Büchi automaton  $A_{\neg\varphi}$ . The third operation is the synchronization between  $\mathcal{K}_M$  and  $A_{\neg\varphi}$ . This constructs a product automaton  $\mathcal{K}_M \otimes A_{\neg\varphi}$  whose language,  $\mathscr{L}(\mathcal{K}_M) \cap \mathscr{L}(A_{\neg\varphi})$ , is the set of executions of M invalidating  $\varphi$ . The last operation is the emptiness check algorithm that explores the product to

tell whether it accepts or not an infinite word, i.e., a counterexample. The model *M* satisfies  $\varphi$  iff  $\mathscr{L}(A_M \otimes A_{\neg \varphi}) = \emptyset$ .

The main problem of model checking is the well known state-space explosion problem. In particular, the performance of the automata-theoretic approach depends (in practice) on the size of the explored part during the emptiness check of the product automaton. This explored part itself depends on three parameters: the automaton  $A_{\neg\varphi}$  obtained from the LTL formula  $\varphi$  to be checked, the Kripke structure  $\mathcal{K}_M$  representing the state-space of the model M, and the emptiness check algorithm. The fact that this algorithm is performed on-the-fly, potentially avoids building the entire product automaton. Indeed, the states of this product that are not visited by the emptiness check are not generated at all.

In order to reduce the size of the product  $\mathcal{K}_M \otimes A_{\neg \varphi}$ , many works have attempted to reduce the size of  $A_{\neg \varphi}$ , either by improving the LTL translation (Section 2.4.6), or by proposing several reductions for the automaton produced by this translation (examples of these reductions are the bisimulation/simulation based reductions presented in Section 2.4.5). However, [78] claims that the size of the product automaton depends more on the "determinism degree" of  $A_{\neg \varphi}$  rather than its size.

Another optimization that can reduce the size of  $A_{\neg \varphi}$  is to simplify the LTL formula  $\neg \varphi$  before translating it into an automaton, several rewriting rules have been proposed [81, 40, 11] to perform LTL simplifications.

According to [27, 49], moving the accepting conditions from the states (as in GBA) to the transitions (TGBA) also reduces the size  $A_{\neg\phi}$ . In addition, any algorithm that translates LTL into a Büchi automaton has to deal with generalized Büchi acceptance conditions at some point, for instance the obtained automaton is a GBA or a TGBA, and the process of degeneralizing this generalized automaton to obtain a BA often increases its size (see Section 2.4.2.1). Several emptiness-check algorithms can deal with generalized Büchi acceptance conditions, making such a degeneralization unnecessary and even costly [29].

It is important to note that in this thesis, we only focus on optimizing the property automaton  $A_{\neg\varphi}$ . We do not consider the techniques whose aim is to reduce the model's state-space  $\mathcal{K}_M$ , such as the *partial order reduction* implemented in Spin tool [55]. Several partial order reduction techniques have been proposed, as the *stubborn sets* of Valmari [86], the *persistent sets* of Godefroid [50] and the *ample sets* of Peled [69]. However, these techniques require additional knowledge about the model. In addition, the state-space of the model is more difficult to manipulate than the property automaton, which is generally smaller.

In the next chapter, we will present the results of an experimentation that compares the performance of different automata-based approaches to model checking, including those using BA versus TGBA, and another kind of  $\omega$ -automaton, called Testing Automata (TA), that is specific to represent stutter-invariant LTL formulas.

# Chapter 3

# **Evaluation of the Testing Automata Approach**

### Contents

3.1	Introduction									
3.2	Stutter-invariant Languages 4									
3.3	Testing Automata (TA): a natural way to monitor the stuttering									
3.4	TA Construction									
	3.4.1	From BA to 0-TA: Construction of an intermediate 0-TA from a Büchi Automaton BA	44							
	3.4.2	From Ø-TA to TA: Elimination of useless stuttering transitions (Ø) and in- troducing livelock-acceptance	45							
	3.4.3	TA Optimizations (that are not yet implemented)	46							
3.5	Explic	it Model checking using TA	47							
	3.5.1	Synchronous Product of a TA with a Kripke structure	47							
	3.5.2	A two-pass emptiness check algorithm	48							
3.6	Exper	imental Comparison of TA versus TGBA and BA	52							
	3.6.1	Implementation on top of Spot	52							
	3.6.2	Benchmark Inputs	53							
	3.6.3	Results	54							
	3.6.4	Discussion (TA two-pass emptiness check problem)	61							
3.7	Concl	usion	64							

# 3.1 Introduction

The previous chapter presented the classical automata-theoretic approach to model checking based on Büchi automata. The main limitation of this approach is the large size of the product automaton  $(\mathcal{K}_{\mathcal{M}} \otimes A_{\neg \varphi})$  obtained by synchronizing the Kripke structure of the model  $\mathcal{K}_{\mathcal{M}}$  with the Büchi automaton  $A_{\neg \varphi}$ , which represents the negation of the LTL property to be checked.

Different variants of Büchi automata have been used with the automata-theoretic approach. In Spot [64, 35], the model checking library we used in the experiments presented in this thesis, LTL properties are represented using TGBA (i.e., the variant of Büchi automata with generalized Büchi acceptance conditions *on transitions* rather on states). Indeed, according to [36], it is preferable

to use TGBA for their conciseness. Unfortunately, having a smaller property automaton  $A_{\neg \varphi}$  does not always imply a smaller product automaton  $\mathcal{K}_M \otimes A_{\neg \varphi}$ . Instead of targeting smaller property automata, some people have attempted to build *more deterministic* [78] ones; however even this does not guarantee the product to be smaller [3, 10].

This chapter focuses on another kind of  $\omega$ -automaton called *Testing Automaton (TA)*. TA are a variant of an "extended" Büchi automata introduced by Hansen et al. [52]. Instead of observing the valuations on states or transitions, the TA transitions only record the changes between these valuations. In addition, TA are less expressive than Büchi automata since they are able to represent only stutter-invariant LTL properties. Also they are often a lot larger than their equivalent Büchi automaton, but their high degree of determinism [52] often leads to a smaller product size [46].

We first provide a detailed presentation of TA and their associated operations for model checking. Then, in order to evaluate the efficiency of LTL model checking using TA, we report and discuss the results of an experimental comparison of three kinds of  $\omega$ -automata: classical Büchi Automata (BA), Transition-based Generalized Büchi automata (TGBA), and Testing Automata (TA) (this part completes our experiment presented in [5]). Our main motivation is to find the technique that seems the most suitable to check a given *stutter-invariant* property on a given model. This is of interest when a tool offers the choice of several techniques, which is the case for the Spot [64] tool in which I implemented the TA approach.

### **3.2** Stutter-invariant Languages

For any  $\omega$ -automaton  $\mathcal{A}$ , we say that a language  $\mathscr{L}(\mathcal{A})$  is *stutter-invariant* if the number of the successive repetitions of any letter of a word  $\sigma \in \mathscr{L}(\mathcal{A})$  does not affect the membership of  $\sigma$  to  $\mathscr{L}(\mathcal{A})$  [39]. In other words,  $\mathscr{L}(\mathcal{A})$  is stutter-invariant *iff* for any finite sequence  $u \in \Sigma^*$ , any element  $\ell \in \Sigma$ , and any infinite sequence  $v \in \Sigma^{\omega}$  we have  $u\ell v \in \mathscr{L}(\mathcal{A}) \iff u\ell\ell v \in \mathscr{L}(\mathcal{A})$ .

Two infinite words  $w_1$  and  $w_2$  are *stuttering equivalent iff* they are equal after removing all repeated letters.

Two languages  $\mathscr{L}(\mathcal{A})$  and  $\mathscr{L}(\mathcal{B})$  are *stuttering equivalent iff* any word of  $\mathscr{L}(\mathcal{A})$  is stuttering equivalent to a word of  $\mathscr{L}(\mathcal{B})$  and vice versa.

Given a stutter-invariant LTL formula  $\varphi$  (Definition 2.3.3) and an  $\omega$ -automaton  $\mathcal{A}_{\varphi}$  such that  $\mathscr{L}(\mathcal{A}_{\varphi}) = \mathscr{L}(\varphi)$ , then  $\mathscr{L}(\mathcal{A}_{\varphi})$  is a stutter-invariant language (we also say that  $\mathcal{A}_{\varphi}$  is a stutter-invariant automaton).

# **3.3** Testing Automata (TA): a natural way to monitor the stuttering

Testing Automata were introduced by Hansen et al. [52] to represent stutter-invariant properties. While a Büchi automaton observes the value of the atomic propositions AP, the basic idea of TA is to only detect the *changes* in these values, making TA particularly suitable for stutterinvariant properties; if a valuation of AP does not change between two consecutive valuations of an execution, the TA can stay in the same state, this kind of transitions are called stuttering transitions. To detect infinite executions that end stuck in the same TA state because they are stuttering,



Figure 3.1: A TA T recognizing the LTL formula  $a \cup Gb$ .

a new kind of accepting states is introduced: *livelock-accepting states*.

If *A* and *B* are two valuations,  $A \oplus B$  denotes the symmetric set difference, i.e., the set of atomic propositions that differ (e.g.,  $a\bar{b} \oplus ab = \{b\}$ ). Technically, this is implemented with an XOR operation (also denoted by the symbol  $\oplus$ ).

**Definition 23** (TA). A TA over the alphabet  $\Sigma = 2^{AP}$  is a tuple  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F}, \mathcal{G} \rangle$ , where:

- *Q* is a finite set of states,
- $I \subseteq Q$  is the set of initial states,
- U: I → 2<sup>Σ</sup> is a function mapping each initial state to a set of valuations (set of possible initial configurations),
- δ ⊆ Q × (Σ \ 0) × Q is the transition relation where each transition (s,k,d) is labeled by a changeset: k ∈ Σ is interpreted as a non empty set of atomic propositions whose value must change between states s and d,
- $\mathcal{F} \subseteq Q$  is a set of Büchi-accepting states,
- $G \subseteq Q$  is a set of livelock-accepting states.

An infinite word  $\sigma = \ell_0 \ell_1 \ell_2 \dots \in \Sigma^{\omega}$  is accepted by  $\mathcal{T}$  iff there exists an infinite sequence  $r = (q_0, \ell_0 \oplus \ell_1, q_1)(q_1, \ell_1 \oplus \ell_2, q_2) \dots (q_i, \ell_i \oplus \ell_{i+1}, q_{i+1}) \dots \in (Q \times \Sigma \times Q)^{\omega}$  such that:

- $q_0 \in I$  with  $\ell_0 \in U(q_0)$ ,
- $\forall i \in \mathbb{N}$ , either  $(q_i, \ell_i \oplus \ell_{i+1}, q_{i+1}) \in \delta$  (the execution progresses in the TA), or  $\ell_i = \ell_{i+1} \land q_i = q_{i+1}$  (the execution is stuttering and the TA does not progress),
- Either, ∀i ∈ N, (∃j ≥ i, ℓ<sub>j</sub> ≠ ℓ<sub>j+1</sub>) ∧ (∃l ≥ i, q<sub>l</sub> ∈ 𝓕) (the TA is progressing in a Büchi-accepting way), or, ∃n ∈ N, (q<sub>n</sub> ∈ 𝔅 ∧ (∀i ≥ n, q<sub>i</sub> = q<sub>n</sub> ∧ ℓ<sub>i</sub> = ℓ<sub>n</sub>)) (the sequence reaches a livelock-accepting state and then stays on that state because the execution is stuttering).

*The language accepted by* T *is the set*  $\mathscr{L}(T) \subseteq \Sigma^{\omega}$  *of executions it accepts.* 

To illustrate this definition, let us consider Figure 3.1, representing a TA  $\mathcal{T}$  for  $a \cup Gb$ . In this figure, the initial states 1, 2 and 3 are labeled respectively by the set of valuations  $U(1) = \{a\bar{b}\}$ ,  $U(2) = \{ab\}$  and  $U(3) = \{\bar{a}b\}$ . Each transition of  $\mathcal{T}$  is labeled with a changeset over the set of atomic propositions  $AP = \{a, b\}$ . In a TA, states with a double enclosure belong to either  $\mathcal{F}$  or  $\mathcal{G}$ : states in  $\mathcal{F} \setminus \mathcal{G}$  have a double plain line, states in  $\mathcal{G} \setminus \mathcal{F}$  have a double dashed line (states 2 and 3 of  $\mathcal{T}$ ), and states in  $\mathcal{F} \cap \mathcal{G}$  use a mixed dashed/plain style (state 4).

- The infinite word *ab*; *āb*; *ab*; *āb*; *ab*; *āb*; *ab*; *a*
- The infinite word *ab*; *āb*; *āb*; *āb*; ... is accepted by a livelock accepting run of *T*. An accepting run starts in state 2, then moves to state 4, and stutters on this livelock-accepting state. Another possible accepting run goes from state 2 to state 3 and stutters in 3 ∈ *G*.
- The infinite word  $ab; a\overline{b}; ab; a\overline{b}; ab; a\overline{b}; \ldots$  is not accepted. It would correspond to a run alternating between states 2 and 1, but such a run is neither Büchi accepting (does not visit any  $\mathcal{F}$  state) nor livelock-accepting (it passes through state  $2 \in \mathcal{G}$ , but does not stay into this state continuously).

### **Property 2.** The language accepted by a testing automaton is stutter-invariant.

*Proof.* This follows from definition of accepted infinite words: a TA may not change its state when an infinite word stutters, so stuttering is always possible.  $\Box$ 

# **3.4 TA Construction**



Figure 3.2: The two steps of the construction of a TA from a BA.

A TA is constructed from a BA in two steps as illustrated in Figure 3.2. The first step constructs an intermediate form of TA, called  $\emptyset$ -TA (for "empty-changesets" TA), which can contain stuttering transitions between two distinct states (i.e.,  $\delta \subseteq Q \times \Sigma \times Q$  in  $\emptyset$ -TA). The second step allows to eliminate these useless stuttering transitions and thus to obtain a TA that satisfies the Definition 23 (i.e.,  $\delta \subseteq Q \times (\Sigma \setminus \emptyset) \times Q$  in TA).

We begin by formally defining an 0-TA automaton and how it accepts infinite words:

**Definition 24** (0-TA). An 0-TA over the alphabet  $\Sigma = 2^{AP}$  is a tuple  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F}, \mathcal{G} \rangle$ , where:

- *Q* is a finite set of states,
- $I \subseteq Q$  is the set of initial states,
- $U: I \to 2^{\Sigma}$  is a function mapping each initial state to a set of valuations,
- δ ⊆ Q × Σ × Q is the transition relation where each transition (s,k,d) is labeled by a changeset k ∈ Σ interpreted as a set of atomic propositions whose value must change between states s and d,
- $\mathcal{F} \subseteq Q$  is a set of Büchi-accepting states,
- $G \subseteq Q$  is a set of livelock-accepting states.





(a) Initial BA for  $a \cup Gb$ .



(c) TA after the simplifications from property 4.

(b) 0-TA after the construction from property 3.



(d) TA after bisimulation.

Figure 3.3: Steps of the construction of a TA from a BA. States with a double enclosure belong to either  $\mathcal{F}$  or G: states in  $\mathcal{F} \setminus \mathcal{G}$  have a double plain line, states in  $\mathcal{G} \setminus \mathcal{F}$  have a double dashed line, and states in  $\mathcal{F} \cap \mathcal{G}$  use a mixed dashed/plain style.

An infinite word  $\sigma = \ell_0 \ell_1 \ell_2 \ldots \in \Sigma^{\omega}$  is accepted by  $\mathcal{T}$  iff there exists an infinite sequence  $r = (q_0, \ell_0 \oplus \ell_1, q_1)(q_1, \ell_1 \oplus \ell_2, q_2) \ldots (q_i, \ell_i \oplus \ell_{i+1}, q_{i+1}) \ldots \in (Q \times \Sigma \times Q)^{\omega}$  such that:

- $q_0 \in I$  with  $\ell_0 \in U(q_0)$ ,
- $\forall i \in \mathbb{N}$ , either  $(q_i, \ell_i \oplus \ell_{i+1}, q_{i+1}) \in \delta$  (we are always progressing in  $\emptyset$ -TA),
- *Either*,  $\forall i \in \mathbb{N}, (\exists j \ge i, q_j \in \mathcal{F})$  (the 0-TA is progressing in a Büchi-accepting way), or,  $\exists n \in \mathbb{N}, \forall i \ge n, (\ell_i = \ell_n) \land (q_i \in G)$  (a suffix of the infinite word stutters in G).

*The language accepted by*  $\mathcal{T}$  *is the set*  $\mathscr{L}(\mathcal{T}) \subseteq \Sigma^{\omega}$  *of infinite words it accepts.* 

Figure 3.3b shows an example of an  $\emptyset$ -TA. The  $\emptyset$ -TA does not respect the constraint of TA that prohibits having stuttering transitions between two distinct states. The second difference is that the stuttering transitions are explicitly represented in an  $\emptyset$ -TA. In other words, an  $\emptyset$ -TA does not use the implicit stuttering of TA (which consists of having an implicit stuttering self-loop on each state of TA).

# 3.4.1 From BA to 0-TA: Construction of an intermediate 0-TA from a Büchi Automaton BA

Geldenhuys and Hansen [46] have shown how to convert a BA into a TA by first converting the BA into an automaton with valuations on the states (called state-labeled Büchi automaton (SLBA)), and then converting this SLBA into an intermediate form of TA (i.e., an  $\emptyset$ -TA) by computing the difference between the labels of the source and destination states of each transition. The next proposition implements these first steps.

**Property 3** (Converting a BA into an 0-TA (an intermediate form of TA) [46]). For any BA  $\mathcal{B} = \langle Q_{\mathcal{B}}, I_{\mathcal{B}}, \delta_{\mathcal{B}}, F_{\mathcal{B}} \rangle$  over the alphabet  $\Sigma = 2^{AP}$  and such that  $\mathcal{L}(\mathcal{B})$  is stutter-invariant, let us define the 0-TA  $\mathcal{T} = \langle Q_{\mathcal{T}}, I_{\mathcal{T}}, U_{\mathcal{T}}, \delta_{\mathcal{T}}, F_{\mathcal{T}}, \emptyset \rangle$  with:

- $Q_T = Q_B \times \Sigma$ ,
- $I_T = I_B \times \Sigma$ ,

• 
$$\forall (s,\ell) \in I_T, U_T((s,\ell)) = \{\ell\},\$$

•  $\forall (s,\ell) \in Q_{\mathcal{T}}, \forall (s',\ell') \in Q_{\mathcal{T}},$  $((s,\ell),\ell \oplus \ell',(s',\ell')) \in \delta_{\mathcal{T}} \iff ((s,\ell,s') \in \delta_{\mathcal{B}}),$ 

```
• F_{\mathcal{T}} = F_{\mathcal{B}} \times \Sigma.
Then \mathscr{L}(\mathcal{B}) = \mathscr{L}(\mathcal{T}).
```

The proof of Property 3 is similar to the proof of Property 8, which will be given in Section 5.3.1.

Figure 3.3b shows the result of applying this construction to the example of Büchi automaton shown for  $a \cup Gb$ . This testing automaton does not yet use livelock-accepting states (the G set). The next property, again from Geldenhuys and Hansen [46], shows how filling G allows to remove all stuttering transitions (i.e., transitions labeled by  $\emptyset$ ) and therefore build the final form of the TA.



(a) Before stuttering transitions ( $\emptyset$ ) reduction. q' is a Büchi accepting state

(b) After reduction,  $q_0$  and  $q_n$  are livelock-accepting states.

Figure 3.4: Elimination of useless stuttering transitions. The states  $q_0$  and  $q_n$  are added to the set of livelock-accepting states G.

# **3.4.2** From Ø-TA to TA: Elimination of useless stuttering transitions (Ø) and introducing livelock-acceptance

The second step consists in filling  $\mathcal{G}$  to simplify  $\mathcal{T}$ , the intuition of this second step is illustrated by Figure 3.4. For that, we compute all strongly connected components using only stuttering transitions (i.e., transitions labeled by  $\emptyset$ ). If such a SCC is not trivial (i.e., it contains a cycle) and contains a Büchi-accepting state, then add all its states to  $\mathcal{G}$ . Then, add to  $\mathcal{G}$  any state that can reach  $\mathcal{G}$  using only stuttering transitions. Finally remove all stuttering transitions from  $\delta$ . The following property formalizes this second step.

**Property 4** (Filling  $\mathcal{G}$  to obtain a TA [46]). Let  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F}, \mathcal{G} \rangle$  be  $\emptyset$ -TA such that  $\mathscr{L}(\mathcal{T})$  is stutter-invariant. By combining the first three of the following operations we can remove all stuttering transitions  $(q, \emptyset, q')$  from  $\mathcal{T}$  and thus obtain a TA. The fourth simplification can be performed along the way.

- If Q' ⊆ Q is a Strongly Connected Component (SCC) such that Q' ∩ 𝔅 ≠ 𝔅 (it is Büchi-accepting), and any two states q,q' ∈ Q' can be connected using a non-empty sequence of stuttering transitions (q,𝔅,q<sub>1</sub>) · (q<sub>1</sub>,𝔅,q<sub>2</sub>) · · · (q<sub>n</sub>,𝔅,q') ∈ δ\*, then the automaton 𝔅' = ⟨Q, I, U, δ, 𝔅, 𝔅, 𝔅 ∪ Q'⟩ is such that 𝔅(𝔅') = 𝔅(𝔅). Such a component Q' is called an accepting Stuttering-SCC.
- 2. If there exists a transition  $(q, \emptyset, q') \in \delta$  such that  $q' \in G$ , then the  $\emptyset$ -TA  $\mathcal{T}'' = \langle Q, I, U, \delta, \mathcal{F}, G \cup \{q\} \rangle$  is such that  $\mathcal{L}(\mathcal{T}'') = \mathcal{L}(\mathcal{T})$ .
- 3. Let T<sup>†</sup> = ⟨Q, I, U, δ, F, G<sup>†</sup>⟩ be the 0-TA obtained after repeating the previous two operations as much as possible (i.e., G<sup>†</sup> contains all states that can be added by the above two operations (Figure 3.4b)). Then, because L(T) and thus L(T<sup>†</sup>) are stutter-invariant, we can remove all stuttering transitions form T<sup>†</sup> to obtain a TA (since stuttering can be captured by the implicit stuttering of TA and the livelock-accepting states of G<sup>†</sup> after the previous two operations). After this last reduction of stuttering transitions, we obtain the final TA T<sup>'''</sup>. Formally, the TA T<sup>'''</sup> = ⟨Q, I, U, δ<sup>'''</sup>, F, G<sup>†</sup>⟩ with δ<sup>'''</sup> = {(q, k, q') ∈ δ | k ≠ 0} is such that L(T<sup>'''</sup>) = L(T<sup>†</sup>) = L(T<sup>†</sup>).
- 4. Any state from which one cannot reach a Büchi-accepting cycle nor a livelock-accepting state can be removed without changing the automaton's language.

The proof of Property4 is similar to the proof of Property 9, which will be given in Section 5.3.2.

The resulting TA can be further simplified by merging bisimilar states (two states are bisimilar if the automaton can accept the same infinite words starting for either of these states). This can be achieved using any algorithm based on partition refinement, the same as for Büchi automata presented in section 2.4.5, taking  $\{\mathcal{F} \cap \mathcal{G}, \mathcal{F} \setminus \mathcal{G}, \mathcal{G} \setminus \mathcal{F}, Q \setminus (\mathcal{F} \cup \mathcal{G})\}$  as initial partition.

Figure 3.3 shows how a BA denoting the LTL formula  $a \cup Gb$  is transformed into a TA by applying prop. 3, prop. 4, and finally merging bisimilar states.

A TA for  $GFa \wedge GFb$  is too big to be shown: even after simplifications it has 11 states and 64 transitions.

### **3.4.3** TA Optimizations (that are not yet implemented)

Looking at Figure 3.3 inspires two optimizations. The first one is based on the fact that the construction of testing automata, described in the previous section, generates a lot of bisimilar states such as the two states labeled with  $(Gb, \bar{a}b)$  and (Gb, ab). This is because the construction considers all the elements of  $\Sigma$  that are compatible with Gb. Had the LTL formula been over  $AP = \{a, b, c\}$ , e.g.,  $(a \lor c) \sqcup Gb$ , then we would have had four bisimilar states:  $(Gb, \bar{a}b\bar{c})$ ,  $(Gb, \bar{a}b\bar{c})$ ,  $(Gb, ab\bar{c})$ , and (Gb, abc). These states are *necessarily* isomorphic, because they only differ in *a* and *c*, some propositions that the formula Gb does not observe.

A more efficient way to construct the testing automaton (and to construct the automaton from Figure 3.3d directly) would be to consider only the subset of atomic propositions that are observed by the corresponding state of the Büchi automaton or its descendants (if the state is labeled by an LTL formula, the atomic propositions occurring in this formula give an over-approximation of that set).

A second optimization relies on the fact any state that does not belong to an SCC can be added to  $\mathcal{F}$  without changing the language of the automaton (this is also true for Büchi automata). For instance on Figure 3.3 the state labeled  $(a \cup Gb, \bar{a}b)$  can be added to  $\mathcal{F}$ . Since this state is not part of any cycle, it cannot occur infinitely often and therefore cannot change the accepted language of the automaton.

$$a\bar{b} \rightarrow (a \cup Gb, a\bar{b}) \{a,b\} \cap \{a\}$$

$$\{b\} \uparrow (b) = \{b\} \circ (Gb,b) = \{a\}$$

$$ab \rightarrow (a \cup Gb, a\bar{b}) \circ (Gb,b) = \{a\}$$

$$ab \rightarrow (a \cup Gb, a\bar{b}) \circ (Gb,b) = \{a\}$$

Figure 3.5: Reduced TA for  $a \cup Gb$ .

This change allows further simplifications by bisimulation: the state  $(a \cup Gb, \bar{a}b)$  is now obviously equivalent to the (Gb,b) state. Figure 3.5 shows the resulting automaton. Note that putting any trivial SCC *x* in  $\mathcal{F}$  before performing bisimulation could hinder the reduction if *x* was isomorphic to some state not in  $\mathcal{F}$ . However if *x* has only successors in  $\mathcal{F}$ , as in our example, then it can be put safely in  $\mathcal{F}$ : indeed, it can only be isomorphic to an  $\mathcal{F}$ -state, or to another

trivial SCC that will be added to  $\mathcal{F}$ . This condition is similar to the one used by Löding before minimizing deterministic weak  $\omega$ -automata [61].

Unfortunately, these optimizations have not been implemented for this work and we had no time to implement the *simulation* [40, 44] reduction.

However, the bisimulation reduction (section 2.4.5) was implemented for the three approaches

BA, TGBA and TA, and used in the experimental comparisons presented later in this thesis.

# 3.5 Explicit Model checking using TA

A first difference between the BA and TA approaches resides in the computation of the synchronous product. Indeed, during the computation of the product between a TA T and a Kripke structure K, T remains in the same state when K executes a stuttering step.

The emptiness check also requires a dedicated algorithm because there are two ways to accept an infinite word: Büchi accepting or livelock accepting. In the algorithm sketched by Geldenhuys and Hansen [46], a first pass is used with an heuristic to detect both Büchi and livelock accepting cycles. Unfortunately, in certain cases, this first pass fails to report existent livelock accepting cycles. This implies that when no counterexample is found by the first pass, a second one is required to double-check for possible livelock accepting cycles. Thus, when there is no counterexample (i.e., the property is satisfied), the entire state-space has to be explored twice.

### 3.5.1 Synchronous Product of a TA with a Kripke structure

For traditional Büchi automata, the product of a BA or a TGBA with a Kripke structure, is also respectively a BA or a TGBA. In the case of testing automata, the product of a Kripke structure and a TA is not a TA. Indeed, while an execution in a TA is allowed to stutter on any state, the execution in a product must execute an explicit stuttering transition [46]. This product automaton can be seen as an  $\emptyset$ -TA but with a small difference in the way of recognizing Büchi-accepting runs. Indeed, a Büchi-accepting run of the product must visit at least a non-stuttering transition infinitely often.

Chapter 5 will introduce a new type of  $\omega$ -automata that improve the TA, called TGTA [6], which simplifies the definition of the product and improves its emptiness check (by proposing a single-pass algorithm). Indeed, the  $\emptyset$ -TGTA used as an intermediate form in the construction of a TGTA will also serve to represent the product between a TGTA and a Kripke structure. In addition, the emptiness check of this product using TGTA will be done in a single pass (in the next section, we will show that the product using TA requires two passes for its emptiness check).

**Definition 25.** For a Kripke structure  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$  and a TA  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F}, \mathcal{G} \rangle$ , the product  $\mathcal{K} \otimes \mathcal{T}$  is a tuple  $\langle S_{\otimes}, I_{\otimes}, U_{\otimes}, \delta_{\otimes}, \mathcal{F}_{\otimes}, \mathcal{G}_{\otimes} \rangle$  where

- $S_{\otimes} = S \times Q$ ,
- $I_{\otimes} = \{(s,q) \in \mathcal{S}_0 \times I \mid l(s) \in U(q)\},$
- $\forall (s,q) \in I_{\otimes}, U_{\otimes}((s,q)) = \{l(s)\},\$
- $\delta_{\otimes} = \{((s,q),k,(s',q')) \mid (s,s') \in \mathcal{R}, (q,k,q') \in \delta, k = l(s) \oplus l(s')\}, \cup \{((s,q), \emptyset, (s',q')) \mid (s,s') \in \mathcal{R}, q = q', l(s) = l(s')\}$
- $\mathcal{F}_{\otimes} = \mathcal{S} \times \mathcal{F}$ ,
- and  $\mathcal{G}_{\otimes} = \mathcal{S} \times \mathcal{G}$ .

An execution  $\sigma = \ell_0 \ell_1 \ell_2 \ldots \in \Sigma^{\omega}$  is accepted by  $\mathcal{K} \otimes \mathcal{T}$  if there exists an infinite sequence  $r = (s_0, \ell_0 \oplus \ell_1, s_1)(s_1, \ell_1 \oplus \ell_2, s_2) \ldots (s_i, \ell_i \oplus \ell_{i+1}, s_{i+1}) \ldots \in (\mathcal{S}_{\otimes} \times \Sigma \times \mathcal{S}_{\otimes})^{\omega}$  such that:

- $s_0 \in \mathcal{S}^0_{\otimes}$  with  $\ell_0 \in U_{\otimes}(s_0)$ ,
- $\forall i \in \mathbb{N}, (s_i, \ell_i \oplus \ell_{i+1}, s_{i+1}) \in \delta_{\otimes}$  (we are always progressing in the product)
- Either,  $\forall i \in \mathbb{N}, (\exists j \ge i, \ell_j \ne \ell_{j+1}) \land (\exists l \ge i, s_l \in \mathcal{F}_{\otimes})$  (the automaton is progressing in a Büchi-accepting way), or,  $\exists n \in \mathbb{N}, \forall i \ge n, (\ell_i = \ell_n) \land (s_i \in \mathcal{G}_{\otimes})$  (a suffix of the execution stutters in  $\mathcal{G}_{\otimes}$ ).

We have  $\mathscr{L}(\mathfrak{K} \otimes \mathcal{T}) = \mathscr{L}(\mathfrak{K}) \cap \mathscr{L}(\mathcal{T})$  by construction.



Figure 3.6: Example of a synchronous product between a Kripke structure  $\mathcal{K}$  and a TA  $\mathcal{T}$  recognizing the LTL formula FG *p*. The bold cycle of  $\mathcal{K} \otimes \mathcal{T}$  is livelock-accepting.

Figure 3.6 shows an example of a synchronous product between a Kripke structure  $\mathcal{K}$  and a TA  $\mathcal{T}$  recognizing the LTL formula FG p. Each state of  $\mathcal{K}$  is numbered and labeled with a valuation of atomic propositions (over  $AP = \{p\}$ ) that hold in this state. In the product  $\mathcal{K} \otimes \mathcal{T}$ , states are labeled with a pairs of the form (s,q) with s is a state of  $\mathcal{K}$  and q of  $\mathcal{T}$ , and the livelock accepting states are denoted by a double dashed circle (as in  $\mathcal{T}$ ). We can notice that this product using TA is smaller than the product of Figure 2.12 using Büchi automata, i.e., the product between the same Kripke structure  $\mathcal{K}$  and the TGBA (or BA) recognizing the LTL formula FG p.

### 3.5.2 A two-pass emptiness check algorithm

In this section, we present a two-pass algorithm for the emptiness check of the synchronous product between a TA and a Kripke structure. While the emptiness check proposed by Geldenhuys and Hansen [46] is based on the Tarjan algorithm [82, 83], the emptiness check proposed in this section, consists of Algorithm 2 (first-pass) and Algorithm 3 (second-pass), which are based on the Dijkstra algorithm for SCCs detection [32, 33].

In model checking approach using TA, the emptiness check requires a dedicated algorithm because according to the Definition 25 (of the different ways to accept a sequence), there are two ways to detect an accepting cycle in the product:

- Büchi accepting: a cycle containing at least one Büchi-accepting state (𝓕) and at least one non-stuttering transition (i.e., a transition (s, k, s') with k ≠ ∅),
- livelock accepting: a cycle composed only by stuttering transitions and livelock accepting states (*G*).
- A straightforward emptiness check would have the following two passes:

1 **Input**: A product  $\mathcal{K} \otimes \mathcal{T} = \langle \mathcal{S}_{\otimes}, I_{\otimes}, U_{\otimes}, \delta_{\otimes}, \mathcal{F}_{\otimes}, \mathcal{G}_{\otimes} \rangle$ **2 Result**:  $\top$  if and only if  $\mathscr{L}(\mathcal{K} \otimes \mathcal{T}) = \emptyset$ **3 Data**: *todo*: stack of  $\langle state \in S_{\otimes}, succ \subseteq \delta_{\otimes} \rangle$ SCC: stack of  $(root \in \mathbb{N}, lk \in 2^{AP}, k \in 2^{AP}, acc \subseteq \mathcal{F}_{\otimes}, rem \subseteq \mathcal{S}_{\otimes})$ *H*: map of  $\mathcal{S}_{\otimes} \mapsto \mathbb{N}$  $max \leftarrow 0$ 4 begin if  $\neg$  first-pass() then return  $\bot$ ; 5 return second-pass() 6 7 first-pass() foreach  $s^0 \in I_{\otimes}$  do 8 DFSpush1 ( $\emptyset$ ,  $s^0$ ) while  $\neg todo.empty()$  do 9 if  $todo.top().succ = \emptyset$  then 10 DFSpop() 11 else 12 pick one  $\langle s, k, d \rangle$  off *todo*.top().*succ***if**  $d \notin H$  **then** 13 DFSpush1(k, d) 14 else if H[d] > 0 then 15 mergel (k, H[d]) if (SCC.top().acc  $\neq \emptyset$ )  $\land$ 16 (*SCC.top*(). $k \neq \emptyset$ ) then return  $\perp$ ; if  $(d \in \mathcal{G}_{\otimes}) \land (SCC.top().k = \emptyset)$  then return  $\bot$ ; 17 return  $\top$ 18 DFSpush1 ( $lk \in 2^{AP}$ ,  $s \in S_{\otimes}$ ) 19  $max \leftarrow max + 1$ 20  $H[s] \leftarrow max$ 21 if  $s \in \mathcal{F}_{\otimes}$  then 22 *SCC*.push( $\langle max, lk, \emptyset, \{s\}, \emptyset \rangle$ ) 23 else 24 *SCC*.push( $\langle max, lk, \emptyset, \emptyset, \emptyset \rangle$ ) 25 *todo*.push( $\langle s, \{\langle q, k, d \rangle \in \delta_{\otimes} | q = s\}$ )) 26 27 mergel ( $lk \in 2^{AP}$ ,  $t \in \mathbb{N}$ )  $acc \leftarrow \emptyset$ 28  $r \leftarrow \emptyset$ 29  $k \leftarrow lk$ 30 while t < SCC.top().root do 31  $acc \leftarrow acc \cup SCC.top().acc$ 32  $k \leftarrow k \cup SCC.top().k \cup SCC.top().lk$ 33  $r \leftarrow r \cup SCC.top().rem$ 34 35 SCC.pop()  $SCC.top().acc \leftarrow SCC.top().acc \cup acc$ 36  $SCC.top().k \leftarrow SCC.top().k \cup k$ 37  $SCC.top().rem \leftarrow SCC.top().rem \cup r$ 38

Algorithm 2: The first-pass of the Emptiness check algorithm for TA products, see Algorithm 3 for the second-pass.

```
1 Data: todo: stack of \langle state \in S_{\otimes}, succ \subseteq S_{\otimes} \rangle
             SCC: stack of (root \in \mathbb{N}, rem \subseteq S_{\otimes})
              H: map of \mathcal{S}_{\otimes} \mapsto \mathbb{N}
             max \leftarrow 0; init \leftarrow I_{\otimes}
 2 second-pass()
 3
          while ¬init.empty() do
                pick one s^0 off init if s^0 \notin H then DFSpush2 (\emptyset, s^0);
 4
 5
                while ¬todo.empty() do
                      if todo.top().succ = \emptyset then
 6
                            DFSpop()
 7
                      else
 8
 9
                            pick one d off todo.top().succ if d \notin H then
                                  DFSpush2(d)
10
                            else if H[d] > 0 then
11
                                  merge2 (H[d]) if (d \in \mathcal{G}_{\otimes}) then return \perp ;
12
          return ⊤
13
14 DFSpush2 (s \in S_{\otimes})
15
          max \leftarrow max + 1
          H[s] \leftarrow max
16
17
          SSCC.push(\langle max, \emptyset \rangle)
          todo.push(\langle s, \{d \in \mathcal{S}_{\otimes} | (s, \emptyset, d) \in \delta_{\otimes}\}))
18
          init \leftarrow init \cup {d \in S_{\otimes} \mid (s,k,d) \in \delta_{\otimes} \land k \neq \emptyset}
19
20 merge2 (t \in \mathbb{N})
          r \leftarrow \emptyset
21
          while t < SCC.top().root do
22
23
                r \leftarrow r \cup SCC.top().rem
                SCC.pop()
24
25
          SCC.top().rem \leftarrow SCC.top().rem \cup r
                   Algorithm 3: The second-pass of the TA emptiness check algorithm.
```

- a first pass to detect Büchi accepting cycles, it corresponds to Algorithm 2 without the test at line 17,
- a second pass presented in Algorithm 3 to detect livelock accepting cycles.

It is not possible to merge these two passes into a single DFS: the first DFS requires the exploration of every transition of the product while the second one must consider only stuttering transitions.

The first-pass of Algorithm 2 detects all Büchi-accepting cycles, and with line 17 included in this algorithm, it detects also some livelock-accepting cycles. Since in certain cases it may fail to report some livelock-accepting cycles, a second pass is required to look for possible livelockaccepting cycles.

This first-pass is based on the TGBA emptiness check algorithm presented in Algorithm 1 (page 34) with the following changes:

- In each item *scc* of the *SCC* stack: the field *scc.acc* contains the Büchi-accepting states detected in *scc*, *scc.lk* is analogous to *la* in Figure 2.13 but it stores the *change-set* labeling the transition coming from the previous SCC, and *scc.k* contains the union of all *change-sets* in *scc* (lines 33 and 37).
- After each merge, *SCC*.top() is checked for Büchi-acceptance (line 16) or livelock-acceptance (line 17) depending on the emptiness of *SCC*.top().*k*.

Figure 3.6 illustrates how the first-pass of Algorithm 2 can fail to detect the livelock accepting cycle in a product  $\mathcal{K} \otimes \mathcal{T}$  as defined in Definition 25. In this example,  $G_{\mathcal{T}} = \{1\}$  therefore (3,1) and (2,1) are livelock-accepting states, and  $C_2 = [(3,1) \rightarrow (2,1) \rightarrow (3,1)]$  is a livelock-accepting cycle.

However, the first-pass may miss this livelock-accepting cycle depending on the order in which it processes the outgoing transitions of (3,1). If the transition  $t_1 = ((3,1), \{p\}, (0,0))$  is processed before  $t_2 = ((3,1), \emptyset, (2,1))$ , then the cycle  $C_1 = [(0,0) \rightarrow (1,0) \rightarrow (2,1) \rightarrow (3,1) \rightarrow (0,0)]$  is detected and the four states are merged in the same SCC before exploring  $t_2$ . After this merge (line 16), this SCC is at the top of the SCC stack. Subsequently, when the DFS explores  $t_2$ , the merge caused by the cycle  $C_2$  does not add any new state to the SCC, and the SCC stack remains unchanged. Therefore, the test line 17 still return false because the union SCC.top().k of all change-sets labeling the transitions of the SCC is not empty (it includes for example  $t_1$ 's label:  $\{p\}$ ). Finally, first-pass algorithm terminates without reporting any accepting cycle, missing  $C_2$ .

On the other side, if the first-pass had processed  $t_2$  before  $t_1$ , it would have merged the states (3,1) and (2,1) in an SCC, and would have detected it to be livelock-accepting.

In general, to report a livelock-accepting cycle, the first-pass computes the union of all changesets of the SCC containing this cycle. However, this union may include non-stuttering transitions belonging to other cycles of the SCC. In this case, the second-pass is required to search for livelock-accepting cycles, ignoring the non-stuttering transitions that may belong to the same SCC.

In the next chapter, we propose a Single-pass Testing Automata STA, which allows to obtain a synchronous product in which such mixing of non-stuttering and stuttering transitions will never occur in SCCs containing livelock-accepting cycles, making the second-pass unnecessary.

The second-pass (Algorithm 3) is a DFS exploring only stuttering transitions (line 18). To report a livelock-accepting cycle, it detects "stuttering-SCCs" and tests if they contain a livelock-accepting state (line 12).

Ignoring the non-stuttering transitions during the DFS, may lead to miss some parts of the product so any destination of a non stuttering transition is stored in *init* for later exploration (line 19).

In the algorithm proposed by Geldenhuys and Hansen [46], the first pass uses a heuristic to detect livelock-accepting cycles when possible. This heuristic detects more livelock-accepting cycles than Algorithm 2. In certain cases this first pass may still fail to report some livelock-accepting cycles. Yet, this heuristic is very efficient: when counterexamples exist, they are usually caught by the first pass, and the second is rarely needed. However, when properties are satisfied, the second pass is always required.

**Note.** It is important to say that in the experiments presented in the sequel, we implement Algorithm 2 including the heuristic proposed by Geldenhuys and Hansen [46] to detect more livelock-accepting cycles during the first pass of the TA approach. We don't present the details of this heuristic because we show in the next chapters other solutions that allow to detect all the livelock-accepting cycles during the first pass and therefore remove the second pass.

# 3.6 Experimental Comparison of TA versus TGBA and BA

This section presents our experimentation of the various types of automata within our tool Spot [64]. We first present the Spot architecture and the way the variation on the model checking algorithm was introduced. Then we present our benchmarks (formulas and models) prior to the description of our experiments.

### 3.6.1 Implementation on top of Spot

Spot is a model-checking library offering several algorithms that can be combined to build a model checker [36].

Figure 3.7 shows the building blocks we used to implement the three approaches:

• the TGBA and BA approaches that share the same synchronous product construction and emptiness check,



Figure 3.7: The experiment's architecture. Two command-line switches control which one of the three approaches is used to verify an LTL formula on a Kripke structure.

• the dedicated algorithms required by the TA approach.

The construction of TGBA and BA already exists in Spot. In addition, we have implemented all algorithms used in the TA approach (TA construction and TA product and its emptiness check). We have also implemented the bisimulation reduction (section 2.4.5) for three approaches BA, TGBA and TA.

In order to evaluate our approach on "realistic" models, we decided to couple the Spot library with the CheckPN tool [36]. CheckPN implements Spot's Kripke structure interface in order to build on-the-fly a state-space of a Petri net. This Kripke structure is then synchronized with an  $\omega$ -automaton (TGBA, BA, or TA), and fed to the suitable emptiness check algorithm. The latter algorithm drives the on-the-fly construction: only the explored part of the product (and the associated states of the Kripke structure) will be constructed.

Constructing the state-space on-the-fly is a double-edged optimization. Firstly, it saves memory, because the state-space is computed as it is explored and thus, does not need be stored. Secondly, it also saves time when a property is violated because the emptiness check can stop as soon as it has found a counterexample. However, on-the-fly exploration is costlier than browsing an explicit graph: an emptiness check algorithm such as the one for TA [52] that does two traversals of the full state-space in the worst case (e.g. when the property holds) will pay twice the price of that construction.

In the CheckPN implementation of the Kripke structure, the Petri Net marking are compressed to save memory. The marking of a state has to be uncompressed every time we compute its successors, or when we compute the value of the atomic properties on this state. These two operations often occur together, so there is a one-entry cache that prevents the marking from being uncompressed twice in a row.

### 3.6.2 Benchmark Inputs

We selected some Petri net models and formulas to compare these approaches.

**Case Studies** The following two bigger models, were taken from actual case studies. They come with some *dedicated* properties to check.

**PolyORB** models the core of the  $\mu$ broker component of a middleware [56] in an implementation using a Leader/Followers policy [71]. It is a Symmetric Net and, since CheckPN processes P/T nets only, it was unfolded into a P/T net. The resulting net, for a configuration involving three sources of data, three simultaneous jobs and two threads (one leader, one follower) is composed of 189 places and 461 transitions. Its state space contains 61 662 states. The authors propose to check that once a job is issued from a source, it must be processed by a thread (no starvation). It corresponds to:

$$\Phi_1 = \mathsf{G}(MSrc_1 \to \mathsf{F}(DOSrc_1)) \land \mathsf{G}(MSrc_2 \to \mathsf{F}(DOSrc_2)) \land \mathsf{G}(MSrc_3 \to \mathsf{F}(DOSrc_3))$$

**MAPK** models a biochemical reaction: Mitogen-activated protein kinase cascade [53]. For a scaling value of 8 (that influences the number of tokens in the initial marking), it contains 22 places and 30 transitions. Its state-space contains  $6.11 \times 10^6$  states. The authors propose to check that from the initial state, it is necessary to pass through states *RafP*, *MEKP*, *MEKPP* and *ERKP*  in order to reach ERKPP. In LTL:

$$\Phi_2 = \neg((\neg RafP) \cup MEKP) \land \neg((\neg MEKP) \cup MEKPP) \land \\ \neg((\neg MEKPP) \cup ERKP) \land \neg((\neg ERKP) \cup ERKPP)$$

Toy Examples A second class of models were selected from the Petri net literature [19, 59]: the flexible manufacturing system (FMS), the Kanban system, the Peterson algorithm, the slotted-ring system, the dining philosophers and the Round-robin mutex [21]. All these models have a parameter n. For the dining philosophers, the Peterson algorithm, the Round-robin, and the slotted-ring, the models are composed of n 1-safe subnets. For FMS and Kanban, n only influences the number of tokens in the initial marking. In our experiments, we selected the following 12 models instances:

- n = 4 and n = 5 for Peterson, FMS and Kanban,
- n = 6 and n = 5 for slotted-ring,
- n = 9 and n = 10 for dining philosophers,
- n = 14 and n = 15 for Round-robin.

**Types of Formulas** As suggested by Geldenhuys and Hansen [46], the type of formula may affect the performances of the various algorithms. In addition to the formulas  $\Phi_1$  and  $\Phi_2$  above, we consider two classes of formulas:

- *RND*: randomly generated LTL formulas (without X operator). Since random formulas are very often trivial to verify (the emptiness check needs to explore only a handful of states), we selected for each model only those formulas requiring more than one second of CPU for the emptiness check in all approaches.
- *WFair*: properties of the form  $(\bigwedge_{i=1}^{n} GFp_i) \rightarrow \varphi$ , where  $\varphi$  is a randomly generated LTL formula. This represents the verification of  $\varphi$  under the weak-fairness hypothesis  $\bigwedge_{i=1}^{n} GFp_i$ . The automaton representing such a formula has at least *n* acceptance conditions which means that the BA will in the worst case be n + 1 times bigger than the TGBA. For the formulas we generated for our experiments we have n = 3.51 on the average.

All formulas were translated into automata using Spot, which was shown experimentally to be very good at this job [74, 34].

For each selected model instance, we generated:

- Verified formulas (i.e., no counterexample in the product): 100 random and 100 weak-fairness,
- Violated formulas (i.e., a counterexample exists): 100 random and 100 weak-fairness.

We consequently have a total 5600 pairs of (model, formula): 2800 violated formulas and 2800 verified formulas.

### 3.6.3 Results

Table 3.1 and Table 3.2 show how BA, TGBA and TA approaches deal with toy models and random formulas.

			Verified properties (no counterexample)									
			1	Automator	n	Full	Emptiness check					
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т		
	BA	avg	13	502	0	2791685	14311851	2791685	14311851	18968		
rson5		max	82	2827	7	11010896	108 553 439	11010896	108553439	80529		
	TGBA	avg	12	460	1	2753024	14166594	2753024	14166594	18656		
ete	TODA	max	54	2816	6	11010896	108 553 439	11010896	108553439	80455		
Ā	TA	avg	83	4324	12	2600539	10293714	5 201 077	20587428	34992		
		max	509	37096	242	9384663	42614845	18769326	85229690	124863		
	D A	avg	13	448	0	1467181	10533598	1467181	10533598	2422		
	DA	max	49	3788	4	4272666	50547840	4272666	50547840	11538		
l go	TCDA	avg	12	411	0	1 398 021	10139043	1398021	10139043	2325		
Ri	TODA	max	49	3788	4	4269384	50547840	4 2 6 9 3 8 4	50547840	10181		
	ТΛ	avg	69	3118	9	1097667	6606130	2195333	13212259	3748		
	іл	max	509	25773	151	2851152	19625864	5702304	39251728	9688		
	Ъ۸	avg	10	267	0	2077187	15 144 149	2077187	15144149	3074		
	DA	max	31	2612	4	9132417	89397363	9132417	89397363	16452		
1S5	тсра	avg	9	245	0	2038054	15037218	2038054	15037218	3017		
E	IUDA	max	28	2612	4	9132417	89397363	9132417	89397363	15594		
	TA	avg	48	1643	5	1401286	11000449	2802572	22000898	5347		
		max	282	11007	100	6109887	54324903	12219774	108649806	26238		
	DA	avg	9	185	0	3494358	33955856	3494358	33955856	5424		
N	DA	max	94	1994	4	22360464	282849140	22360464	282849140	43179		
bar	TCRA	avg	8	166	0	3356053	32800737	3356053	32800737	5253		
Xan	TODA	max	60	1994	4	20253072	258315134	20253072	258315134	44 28 1		
1	ТΔ	avg	44	1659	3	2433853	23389805	4867707	46779609	9192		
	17.1	max	277	21251	82	15272712	161 364 553	30545424	322729106	61873		
	Ъ۸	avg	16	705	1	4952039	29080163	4952039	29080163	9841		
	DA	max	81	5995	7	18399098	121 269 824	18399098	121269824	43970		
101	TGRA	avg	14	636	1	4668178	27564474	4668178	27564474	9352		
Phi	TODA	max	81	5397	6	17947837	119545256	17947837	119545256	40273		
.	ТΔ	avg	71	4092	20	2334154	19893200	4668308	39786400	14494		
	17.1	max	412	55321	232	8378151	74240975	16756302	148481950	53359		
	B٨	avg	16	717	1	2776200	31 544 574	2776200	31544574	9345		
N	DA	max	100	7680	4	15198075	238617470	15198075	238617470	71576		
lii1	TGPA	avg	15	664	0	2682881	30981263	2682881	30981263	9192		
Sob	TODA	max	82	5792	4	13129644	225905117	13129644	225905117	62328		
	ТΔ	avg	95	6048	15	2027251	17168855	4054503	34337710	12401		
	IA	max	481	43547	143	8169472	68645888	16338944	137291776	48675		

Table 3.1: Comparison of the three approaches on toy examples with **random formulæ**, when **counterexamples do not exist.** 

			Violated properties (a counterexample exists)									
				Automator	1	Full	product	Emptiness check				
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т		
	BA	avg	17	738	1	8591425	45706072	608443	2310272	3932		
n5		max	70	3568	5	50366303	365 239 596	3 599 887	21728436	24720		
[OS]	TGRA	avg	14	646	0	8118953	42865329	600 828	2325266	3892		
ete	TODA	max	47	3 3 9 2	5	50297469	364 396 821	3632979	22095045	24665		
Pe	TA	avg	90	5683	15	8 592 222	33972863	546562	1886028	3667		
		max	656	62947	132	56166663	224108070	3178506	11478152	21823		
	DA	avg	20	832	1	3 3 7 8 6 8 4	28026221	876552	5897370	1 3 6 9		
	DA	max	64	4460	3	11012779	123 266 244	6792681	70374257	15784		
lgo	TGRA	avg	17	713	1	3 1 6 3 7 6 3	26260802	854229	5805080	1 3 2 9		
Ri	TODA	max	51	3788	4	9855589	108310156	5914415	65020741	13999		
	ТА	avg	125	7462	14	2846468	18131253	650 592	$\mathbf{3881701}$	1170		
	IA	max	472	41068	162	8960960	62840996	4923301	36659892	9895		
	БА	avg	16	553	1	10949858	101710172	1765510	11843265	2319		
	DA	max	104	4524	3	42512695	432252072	14455789	120456128	25056		
1S5	TGBA	avg	14	496	1	10174866	95061682	1602719	10636471	2148		
F		max	96	4200	3	40672843	404 072 123	11075871	110774947	18644		
	TA	avg	90	4112	9	8 418 689	70784299	1051522	7742600	1956		
		max	390	22702	79	31864749	267656316	8518894	63178516	16558		
	BA	avg	15	391	1	11646608	127363073	1 204 1 37	9094855	1518		
N		max	87	2696	4	43038083	587 543 197	13023468	185410581	25805		
bar	TCDA	avg	13	355	0	10952627	119272618	1153850	8728297	1473		
an	IUDA	max	87	2184	3	43038083	587 543 197	13023468	185410581	26750		
X	ТА	avg	75	2870	6	8 4 2 9 9 4 8	83120910	670 472	5873480	1224		
		max	473	30375	92	31602068	331 833 407	7673335	77 364 100	14239		
	ДΛ	avg	14	641	0	20491935	206666412	1551357	7673496	2713		
	DA	max	74	5928	4	84068722	1 377 479 362	9032250	53881112	17065		
101	TGRA	avg	13	590	0	19118012	195 258 228	1518296	7 522 623	2651		
Phi	TODA	max	74	5928	3	84068722	1 377 479 362	9032250	53881112	15880		
	ТА	avg	92	5690	14	16012398	148019289	900 840	6499314	2545		
		max	356	46498	170	66535322	696331784	7342016	64470840	25536		
	ВЛ	avg	24	1113	1	5866595	71 526 501	1719685	18722554	5 5 4 9		
N	DA	max	84	5928	6	17 192 350	282 229 101	13627374	191356248	53159		
lin1	TGRA	avg	21	965	1	5514573	67871153	1676314	18343280	5475		
Sob		max	74	5928	7	15127431	281 880 971	13627374	191356248	52962		
	Т۸	avg	164	11366	21	4869334	41343229	1 249 161	10420582	4052		
	IA	max	854	62419	107	11791872	104488704	7767216	74241202	28352		

Table 3.2: Comparison of the three approaches on toy examples with **random formulæ**, when **counterexamples exist**.

				)							
			A	Automaton		Full	product	Emptiness check			
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т	
	ЪΛ	avg	9	93	0	3 5 3 0 2 0 2	17259383	3 5 3 0 2 0 2	17259383	23976	
<b>D5</b>	DA	max	32	494	2	14572104	105 106 990	14572104	105106990	99998	
rson	TGRA	avg	4	52	0	3071854	15745483	3071854	15745483	21163	
ete	TODA	max	17	357	2	9551018	87453872	9551018	87453872	70107	
ď	TA	avg	50	564	1	3018533	11771410	6037065	23542820	41176	
		max	184	3362	3	10750205	42228050	21500410	84456100	146637	
	DA	avg	10	104	0	1565890	9845113	1 565 890	9845113	2345	
	DA	max	34	494	2	4945096	44945800	4945096	44945800	9306	
l go	TCDA	avg	5	64	0	1 1 2 1 5 9 0	7920270	1121590	7920270	1903	
Ri	TODA	max	25	357	2	3873576	38047692	3873576	38047692	7910	
	ТЛ	avg	58	682	1	1118511	6494556	2237021	12989111	3697	
	IA	max	183	3353	3	3 3 8 3 3 7 2	22481284	6766744	44962568	10865	
	Ъ٨	avg	9	84	1	3966306	28 568 464	3966306	28568464	5 5 9 6	
	DA	max	25	340	2	14865917	130390060	14865917	130390060	24590	
1S5	тсра	avg	4	41	0	3030750	23 589 355	3030750	23589355	4674	
E	TODA	max	12	212	2	11882973	114787553	11882973	114787553	23008	
	TA	avg	48	515	1	2259230	17600147	4518460	35200294	8482	
		max	171	4256	3	9764223	79707273	19528446	159414546	36382	
	BA	avg	8	62	0	3784743	33114074	3784743	33114074	5190	
2		max	32	221	2	15768592	198359739	15768592	198359739	29290	
bar	TCRA	avg	4	33	1	3098615	29435137	3098615	29435137	4712	
(an	TODA	max	21	168	2	13014212	187004177	13014212	187004177	26143	
1	ТΔ	avg	36	306	1	2083777	19156589	4167555	38313178	7883	
	171	max	140	2049	2	11039112	106534190	22078224	213068380	48724	
	ВЛ	avg	9	80	0	5478383	33455338	5478383	33455338	11245	
	DA	max	31	547	1	16695369	140088618	16695369	140088618	41 598	
101	TGBA	avg	4	35	0	4255979	26793678	4255979	26793678	9064	
Phi	TODA	max	21	216	2	12974557	112561242	12974557	112561242	39806	
	ТА	avg	47	501	1	3229690	27631036	6459379	55262071	20439	
	171	max	289	8199	6	10987384	94141317	21974768	188282634	64593	
	RΔ	avg	9	98	0	2262431	20634259	2262431	20634259	6703	
2	D/1	max	38	608	1	5849069	76754676	5849069	76754676	20953	
ii	TGR∆	avg	5	59	0	1782133	17996764	1782133	17996764	5749	
Rot		max	16	337	1	5029881	70344648	5029881	70344648	19263	
	ΤΔ	avg	54	618	1	1508082	12171559	3016163	24343119	9099	
	1/1	max	239	3706	4	3944448	32391168	7888896	64782336	23018	

Table 3.3: Comparison of the three approaches on toy examples with **weak-fairness formulæ**, **when counterexamples do not exist.** 

			Violated properties (a counterexample exists)									
			1	Automator	ı	Full p	product	Emptiness check				
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т		
	RΔ	avg	12	184	1	9604699	42508151	654497	2230146	4265		
n5	DA	max	74	1338	2	34771276	235328943	7043181	27266804	46628		
LSO	TGBA	avg	5	78	0	3731947	17751951	522414	1759041	3438		
ete	IUDA	max	18	351	2	13508730	79589415	3696481	16733420	24769		
	ТА	avg	95	1400	1	10950499	42790086	641997	2123401	4368		
		max	350	5968	8	30354888	120985972	7637347	29621420	51378		
	BA	avg	18	291	1	4076877	30760035	1312236	7496631	1822		
	DA	max	70	2248	3	6924114	96086800	3706323	34372404	6998		
ng(	TGBA	avg	7	123	1	2000787	15868857	890 567	5736568	1410		
Ri	TODA	max	38	1168	3	4670124	53246768	2766283	26917302	6101		
	ТА	avg	128	2009	2	3456785	21069783	905280	5156350	1583		
	171	max	388	12783	13	6152232	43283696	2747578	17464827	4995		
	BA	avg	12	139	0	13634313	123924884	1623286	8873430	1807		
	DA	max	46	550	2	38283058	436545323	11138596	82531874	14207		
IS	TGBA	avg	5	68	0	7978129	75701408	1 372 877	7 544 397	1607		
E		max	23	370	2	27 570 312	293219919	6676584	53451994	9961		
	TA	avg	71	871	1	10007988	81736556	845087	5687307	1471		
		max	245	3766	5	27733464	249 573 778	6425987	51962113	11856		
	BA	avg	12	133	1	11 528 567	115461763	1541122	10510834	1738		
S.		max	34	501	2	34735366	497454792	12460358	113423964	16932		
lbai	TGBA	avg	5	66	0	7040236	73867374	1 3 4 4 8 7 0	9279129	1562		
Kan	TODA	max	17	305	2	25423161	388961954	7950633	80041971	11336		
-	ТА	avg	69	849	1	8102548	77515606	809 384	6986502	1440		
		max	237	3655	5	32732392	329125090	5 501 242	53375472	10359		
	BA	avg	12	148	0	20724516	190232774	1144023	5417407	1849		
0	DA	max	44	688	3	52231021	996567139	7775968	45330784	13909		
101	TGBA	avg	5	69	0	10515366	102622810	833663	3762217	1357		
Phi		max	25	392	3	24496745	402197001	5017174	32273634	9802		
	ТА	avg	81	1027	1	16779360	150030607	743065	5019598	1936		
	171	max	267	4339	6	43186049	394172486	6703064	55546036	21005		
	RΔ	avg	21	356	1	7007406	78083493	1916523	16640621	5468		
S		max	73	2248	2	23830436	327 588 082	8911386	101 808 220	28268		
l'ii	TGB∆	avg	9	157	1	4006015	46910526	1486839	14360747	4691		
Rot		max	44	1168	2	14802883	195152925	5666009	78112291	21619		
	ТА	avg	143	2467	2	5832674	48201707	1 347 956	10786002	4257		
		max	466	12783	12	18204684	151 155 748	6486142	53115008	18506		

Table 3.4: Comparison of the three approaches on toy examples with **weak-fairness formulæ**, **when counterexamples exist.** 

				Verified properties (no counterexample)									
				1	Automator	1	Full	product	Emptiness check				
				st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т		
		Ъ۸	avg	18	781	1	111714	337 573	111714	337 573	612		
		DA	max	136	7680	5	373 039	2138028	373 039	2138028	2097		
	5 D	тера	avg	16	715	1	109158	328976	109158	328976	601		
	R	TUDA	max	136	5792	4	373 039	2138028	373 039	2138028	2207		
		T۸	avg	94	5768	17	94283	$\mathbf{205464}$	188566	410928	1037		
3/2		IA	max	481	43 547	196	261318	604061	522636	1 208 1 22	2847		
33		DA	avg	11	157	0	132102	369787	132102	369787	713		
R		BA	max	71	1 2 9 4	2	799355	2630204	799355	2630204	4339		
I <sub>V</sub> C	air	TGBA	avg	5	72	0	96435	277426	96435	277426	<b>524</b>		
Po	W		max	18	337	2	368 898	1450005	368 898	1450005	2049		
			avg	79	1106	1	120479	262456	240958	524913	1302		
		IA	max	351	5920	9	633 208	1408943	1266416	2817886	6806		
	$\Phi_1$	BA	-	7	576	1	345 241	760491	345241	760491	1675		
		TGBA	_	7	576	1	345 241	760491	345241	760491	1688		
		TA	-	80	14590	8	342613	742815	685226	1485630	3391		
		BA	avg	12	374	1	2894799	25769242	2894799	25769242	5477		
			max	79	3813	5	13636352	147 555 158	13636352	147 555 158	29292		
	Ð		avg	10	333	1	2808973	25212509	2808973	25212509	5325		
	R	IGBA	max	64	3813	5	13636352	147 555 158	13636352	147 555 158	28979		
		<b>T</b> 4	avg	50	2239	9	1724291	20203618	3448582	40407237	9591		
		IA	max	336	22614	146	8469258	108847708	16938516	217695416	53824		
8		DA	avg	10	82	0	4290714	38013034	4290714	38013034	7887		
<b>I</b> ∎		ВА	max	46	437	2	17600440	177748756	17600440	177748756	33063		
Ž	air	тсра	avg	5	39	0	3898645	34822961	3898645	34822961	7344		
	W	IGBA	max	21	198	2	14452198	162156912	14452198	162156912	31117		
		<b>T</b> 4	avg	44	407	1	2112810	24492299	4225619	48984598	12092		
		IA	max	171	1920	3	6110748	75624744	12221496	151249488	40347		
		BA	-	6	165	0	46494	302350	46 494	302350	45		
	$\Phi_2$	TGBA		6	165	1	46494	302350	46 494	302350	50		
		TA	-	9	293	2	33 376	289235	66752	578470	95		

Table 3.5: Comparison of the three approaches for the **case studies when counterexamples do not exist.** 

			[	Violated properties (a counterexample exists)									
				Automaton			Full	product	Emptiness check				
				st.	tr.	$T_{\boldsymbol{\phi}}$	st.	tr.	st.	tr.	Т		
		B٨	avg	15	839	1	137958	397 495	54412	130427	294		
		DA	max	74	7004	11	1 5 2 0 6 4 9	5898222	192692	514905	1088		
	R	TCRA	avg	14	760	1	132232	380305	53901	128960	291		
	R	TODA	max	65	7004	9	1295661	5045814	192692	514905	1049		
		т۸	avg	81	6223	29	135237	296081	56518	126300	310		
3/2		IA	max	540	59971	631	1 377 784	2976709	243950	547459	1 3 2 4		
33		B٨	avg	11	150	0	190987	472205	112918	279490	600		
R		DA	max	43	788	2	604611	1951062	328442	1419816	1794		
IVC I	Fair	TGBA	avg	5	66	0	96 849	258704	69053	177391	372		
$\mathbf{P}_{0}$	[M]		max	22	516	2	432240	1584372	193264	1 031 898	1110		
		TA	avg	82	1157	1	201 879	434747	113705	250573	612		
			max	262	4636	6	560685	1 194 198	254328	594751	1330		
		BA	avg	15	578	1	29980338	454492066	1939885	13406649	2890		
			max	61	4398	2	141058746	2921365220	13284206	120933281	24053		
	P	тсра	avg	13	513	1	27266223	405 363 157	1850695	12575125	2702		
	R	IUDA	max	59	4298	2	91059214	1963331216	13284206	111916150	21771		
		Т۸	avg	88	4818	13	22458466	298950743	967363	9259668	2346		
		IA	max	334	35401	164	90049281	1 300 904 178	5241327	61655512	13808		
K 8		RΔ	avg	13	176	0	34857814	493 530 217	2249726	16186296	3515		
AP		DA	max	46	541	2	111139060	1700666096	15349474	123 057 668	26758		
Σ	Fair	TCRA	avg	6	81	0	20904101	300416331	1962765	13896236	3129		
	M	TODA	max	23	414	2	101813441	1681905871	9463562	72266873	16928		
		ТΔ	avg	88	1212	1	27535941	358 622 897	1112324	11462361	2869		
		IA	max	245	3749	4	93441682	1258076980	5203746	59642346	14489		

Table 3.6: Comparison of the three approaches for the case studies when counterexamples exist.

Table 3.3 and Table 3.4 show toy models against weak-fairness formulas.

Table 3.5 and Table 3.6 show the results of the two cases studies against random, weak-fairness, and dedicated formulas issued from the studies.

These tables separate cases where formulas are verified from cases where they are violated. In the former (Table 3.1, 3.3 and 3.5), no counterexample are found and the full state-space had to be explored; in the latter (Table 3.2, 3.4 and 3.6) the on-the-fly exploration of the state-space stops as soon as the existence of a counterexample is computed.

All values shown in all tables are averaged over 100 different formulas (except for the lines  $\Phi_1$  and  $\Phi_2$  in Table 3.5, where only one formula is used). For instance we checked Peterson5 against 100 random formulas that had no counterexample, and against 100 random formulas that had a counterexample. The average and maximum are computed separately on these two sets of formulas.

Column-wise, these tables show the average and maximum sizes (states and transitions) of: (1) the automata  $A_{\neg \varphi_i}$  expressing the properties  $\varphi_i$ ; (2) the products  $A_{\neg \varphi_i} \otimes \mathcal{K}$  of the property with the model; and (3) the subset of this product that was actually explored by the emptiness check. For verified properties, the emptiness check of TGBA and BA always explores the full product so these sizes are equal, while the emptiness check of TA always performs two passes on the full product so it shows double values. On violated properties, the emptiness check aborts as soon as it finds a counterexample, so the explored size is usually significantly smaller than the full product.

The emptiness check values show a third column labeled "T": this is the time in  $\frac{1}{100^e}$  of seconds (a.k.a. centiseconds) spent doing that emptiness check, including the on-the-fly computation of the subset of the product that is explored. The time spent constructing the property automata  $A_{\neg \varphi_i}$  is shown in column "T $_{\varphi}$ " (it is negligible compared to that of the emptiness check "T").

Figure 3.8 compares the number of visited transitions when running the emptiness check; plotting TA against BA and TGBA. This gives an idea of their relative performance. Each point corresponds to one of the 5600 evaluated formulas (2800 violated with counterexample as black circles, and 2800 verified having no counterexample as green crosses). Each point below the diagonal is in favor of TA while others are in favor of the other approach. Axes are displayed using a logarithmic scale.

Figure 3.9 compares the number of visited transitions between BA and TGBA. Each point below the diagonal is in favor of TGBA (this clearly shows that BA are less efficient than TGBA).

All these experiments were run on a 64bit Linux system running on an Intel(R) 64-bit Xeon(R) @2.00GHz, with 10GB of RAM.

### 3.6.4 Discussion (TA two-pass emptiness check problem)

Although the state-space of cases studies can be very different from those of random state-spaces [68], a first look at our results confirms two facts already observed by Geldenhuys and Hansen using random state-spaces [46]: (1) although the TA constructed from properties are usually a lot larger than BA, the average size of the full product is smaller thanks to the more deterministic nature of the TA. (2) For violated properties, the TA approach explores less states and transitions on the average than the BA.



(a) TA against BA approaches



(b) TA against TGBA approaches

Figure 3.8: Performance (number of transitions explored by the emptiness check) of TA against BA and TGBA.



(a) TGBA against BA approaches

Figure 3.9: Performance (number of transitions explored by the emptiness check) of TGBA against BA.

We complete this picture by showing execution times, by separating verified properties from violated properties, and by also evaluating the TGBA approach.

**On verified properties**, the results are very straightforward to interpret: the BA are slightly worse than the TGBA because the degeneralization increases the size of the property automaton. In fact, the average number of acceptance conditions needed in random formulas (Table 3.1 and 3.5) is so close to 1 that the degeneralization barely changes the sizes of the automata. With weak-fairness formulas (Table 3.3 and 3.5), the number of acceptance conditions is greater, so TGBA are favored over BA. Surprisingly, both TGBA and BA, although they are not tailored to *stutter-invariant* properties like TA, appear more effective to prove that a *stutter-invariant* property is verified. In the three tables, although the full product of the TA approach is smaller than the other approaches, it has to be explored twice (as explained in section 3.4): the emptiness-check consequently explores more states and transitions. This double exploration is not enough to explain the big runtime differences. Two other subtle implementation details of the synchronous products contribute to the time difference:

To synchronize a transition of a Kripke structure with a transition (or a state in case of stuttering) of a TA, we must compute the symmetric difference (i.e., the changeset) *l(s)* ⊕ *l(d)* between the labels of the source and destination states. The same synchronization in the TGBA and BA approaches requires to know only the source label.

Computing these labels is a costly operation in CheckPN because Petri net marking are compressed in memory to save space. Although we implemented some (limited) cache to

alleviate the number of such label computation, profiling measures revealed the TA approach was 3 times slower than the TGBA and BA approaches, but that labels where computed 9 times more.

 A second implementation difference is that the transitions of testing automata are labeled by elements of Σ, while in the Spot implementation of TGBA and BA, the transitions are labeled by elements of 2<sup>Σ</sup>. That means that once l(s) ⊕ l(d) ∈ Σ has been computed, we can use a hash table to immediately find matching transitions of the testing automaton. In the TGBA and BA implementations, we linearly scan the list of transitions of the property automaton until we find one compatible with l(s).

In order to protect the results against the influence of various optimizations, implementation tricks, and the central processor and memory architecture, Geldenhuys and Hansen [46] have decided (in their work about TA) to report only the number of explored states and transitions, and not the number of bytes and milliseconds consumed by their implementations. They found that the number of states gives a reliable indication of the memory required, and, similarly, the number of transitions a reliable indication of the time consumption. That is why in our discussion (in this chapter and in the next chapters), we focus on the number of explored states and transitions more than time consumption.

**On violated properties**, it is harder to interpret these tables because the emptiness check returns as soon as it finds a counterexample. Changing the order in which non-deterministic transitions of the property automaton are iterated is enough to change the number of states and transitions to be explored before a counterexample is found: in the best case the transition order leads the emptiness check straight to an accepting cycle; in the worst case, the algorithm explores the whole product until it finally finds an accepting cycle. Although the emptiness check algorithms for the three approaches share the same routines to explore the automaton, they are all applied to different kinds of property automata, and thus provide different transition orders.

This ordering luckiness explains why the BA approach sometimes outperforms the TGBA one.

We believe that the TA, since they are more deterministic, are less sensible to this ordering. They also explore a smaller state-space on the average. This smaller exploration is not always tied to a good runtime because of the extra computation of labels discussed in this section. Again, looking at the average number of transitions explored by the emptiness check indicates that the TA approach would outperform the others if labels computation would have been cheap.

# 3.7 Conclusion

Geldenhuys and Hansen evaluated the performance of the BA and TA approaches with small random Kripke structures checked against LTL formulas taken from the literature [46]. In this work, we have completed their experiments by using actual models and different kinds of formulas (random formulas not trivially verifiable, random formulas expressing weak-fairness formulas, and a couple of real formulas), by evaluating the TGBA approach, and by distinguishing violated formulas and verified formulas in the benchmark.

In TA approach, an unfortunate consequence of having two different ways of accepting infinite words (livelock or Büchi), is that the emptiness-check algorithm required during model checking
must perform two passes on the whole state-space in the worst case.

For verified formulas, we found that the state-space reduction achieved by the TA approach was not enough to compensate for the two-pass emptiness check this approach requires. It is therefore better to use the TGBA approach to prove that a *stutter-invariant* formula is verified and TA approach in an earlier "debugging phase".

When the formulas are violated, the TA approach usually processes less transitions than the BA approach and TGBA to find a counterexample. This approach should therefore be a valuable help to debug models (i.e. when counterexamples are *expected*). This is especially true on random formulas. With weak-fairness formulas, generalized automata are advantaged and are able to beat the TA on the average in 4 of our 8 examples (Peterson5, Philo10, Ring6, PolyORB 3/2/2).

In the next chapter, we propose some optimizations in order to omit the second pass in TA approach, in particular when no livelock-accepting states is encountered during the first pass. We also propose STA (*Single-pass Testing Automata*), a transformation of TA that never requires such a second pass.

In Chapter 5, we propose a single-pass and generalized testing automata, called TGTA (*Transition-based Generalized Testing Automata*). TGTA combines ideas from TA and TGBA. The basic idea is to have a form of testing automata with transition-based generalized acceptance conditions, which allows us to modify the automata construction in order to remove the second pass of the emptiness check of the product.

We also noted that making a product between a TA and a Kripke structure is not a good idea when computing the label of the Kripke structure is expensive. A more efficient model checker using testing automata should probably represent the model using labeled Kripke structure in which the transition labels represent the symmetric difference (i.e., the changeset) between the label of the source and destination state. This optimization is exploited in the chapter 6 concerning the symbolic model checking using TGTA.

## Part II

## Improving Testing Automata for Explicit LTL Model Checking

### CHAPTER 4

## Improving the Testing Automata Approach

#### Contents

4.1	1 Introduction											
4.2	Improving the Emptiness check by avoiding the second pass in particular cases 7											
4.3	Converting a TA into a Single-pass Testing Automaton (STA)											
	4.3.1	Single-pass Testing Automata (STA)	71									
	4.3.2	Construction of an STA from a TA	71									
	4.3.3	Correctness of the one-pass emptiness check using STA	72									
	4.3.4	STA optimization	73									
4.4	Exper	imental evaluation of the TA improved emptiness check and of STA	75									
	4.4.1	Implementation	75									
	4.4.2	Results	76									
	4.4.3	Discussion	84									
4.5	Concl	usion	85									

### 4.1 Introduction

In the previous chapter, we evaluated the use of Testing Automata (TA) for the model checking of stutter-invariant properties. We have shown that the TA approach is efficient when the formula to be verified is violated (i.e., a counterexample exists). This is not the case when the property is verified since the entire state-space has to be visited twice to check for each acceptance mode of a TA (Büchi-acceptance or livelock-acceptance).

In this chapter we improve the TA approach in two ways. First, we propose optimizations of the emptiness check algorithm that avoid the second pass when it is possible. Second, we propose a transformation of TA into a normal form that never requires such a second pass, called *Single-pass Testing Automata* (STA).

Although STA are more constrained than TA, we provide a transformation that automatically translates the latter into the former. Then we prove the correctness of the single pass emptiness check for STA. We also present an optimization that allows to build a smaller STA.

We have implemented the algorithms of STA approach in Spot, our model checking library. We are thus able to compare them with the "traditional" algorithms we used on Testing Automata (TA) and Transition-based Generalized Büchi Automata (TGBA). These experiments show that STA compete well on our examples.

# 4.2 Improving the Emptiness check by avoiding the second pass in particular cases

In the evaluation of TA approach presented in the previous chapter, we have implemented the heuristic for livelock detection proposed by Geldenhuys and Hansen [46] to avoid the second pass of the emptiness check algorithm. Unfortunately, this heuristic fails in certain cases to detect livelock cycles. In addition to this heuristic, we add the following improvements to the first pass in order to avoid the second pass in the following particular cases:

- 1. In the previous chapter, we have shown that the second pass is only used to detect livelock-accepting cycles. Therefore, if no livelock-accepting state is visited during the first pass (i.e., the product does not contain livelock-accepting states), then the second pass can be disabled: this can be done by simply adding a variable *Gseen* in Algorithm 2 (page 49), where *Gseen* is a flag that records if a livelock-accepting state is detected during the exploration of the product by the first pass. In the experiments presented at the end of this chapter, this optimization greatly improves the performance of the TA approach in the cases where the formula is verified.
- 2. A cycle detected in the product during the first pass is also accepted if it contains a livelockaccepting state (s,q) of the product such that q has no successor in the TA. Indeed, from this state, a run of the product can only execute stuttering transitions. Therefore, a cycle containing this state, is composed only by stuttering transitions: it is a livelock-accepting cycle.

This second optimization may detect some livelock-accepting cycles, but it misses livelock-accepting cycles that are mixed with non-stuttering transitions in the same SCC, as discussed in section 3.5.2 (page 48) of the previous chapter.

We now introduce a class of TA where the second optimization will always detect any livelockaccepting cycle during the first-pass.

## 4.3 Converting a TA into a Single-pass Testing Automaton (STA)

In this section, we introduce STA, a transformation of TA into a normal form such that livelockaccepting states have no successors, and therefore STA approach does not need the second pass of the emptiness check of TA approach. This contribution improves the efficiency of the model checking (this will be experimentally evaluated in section 4.4). In addition, STA simplify the implementation (and the optimization) of the emptiness check algorithm as it renders unnecessary the implementation of the second pass.

#### 4.3.1 Single-pass Testing Automata (STA)

**Definition 26** (STA). A Single-pass Testing Automaton (STA) is a TA  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F}, \mathcal{G} \rangle$  over  $\Sigma$  such that  $\delta \cap (\mathcal{G} \times 2^{\Sigma} \times Q) = \emptyset$ . In other words, an STA is a TA in which every livelock-accepting state has no successors.

#### 4.3.2 Construction of an STA from a TA



Figure 4.1: Example of a product between a Kripke structure  $\mathcal{K}$  and a TA  $\mathcal{T}$  of FG p. The bold cycle of  $\mathcal{K} \otimes \mathcal{T}$  is livelock-accepting.



Figure 4.2: Impact on the product of using STA  $T^+$  instead of TA T. Bold states and transitions are addition relative to Figure 4.1.

Property 5 formalizes the construction of an STA from a TA. We can transform a TA into an STA by adding an unique livelock-accepting state g (i.e., in STA,  $G = \{g\}$ ), and adding a transition (q,k,g) for any transition (q,k,q') that goes into a livelock-accepting state  $q' \in G$  of the original automaton. In addition, if q' has no successors then q' can be removed, since it is bisimilar to the new state g.

**Property 5.** Let  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F}, \mathcal{G} \rangle$  be a TA, the equivalent STA is  $\mathcal{T}' = \langle Q', I', U', \delta_{\oplus}', \mathcal{F}, \{g\} \rangle$  where

- $Q' = (Q \setminus G_0) \cup \{g\}$  where  $G_0 = \{q \in G \mid (\{q\} \times \Sigma \times Q) \cap \delta = 0\}$  is the set of states of G that have no successors, and  $g \notin Q$  is a new state,
- $I' = I \cup \{g\}$  if  $\mathcal{G} \cap I \neq \emptyset$ , I' = I otherwise,
- $\delta'_{\oplus} = (\delta \setminus (Q \times \Sigma \times G_{\emptyset})) \cup \{(q,k,g) \mid (q,k,q') \in \delta, q' \in \mathcal{G}\},\$

•  $\forall q \in I, U'(q) = U(q) \text{ and } U'(g) = \bigcup_{q \in (\mathcal{G} \cap I)} U(q).$ 

is such that  $\mathscr{L}(\mathcal{T}') = \mathscr{L}(\mathcal{T})$ .

Figure 4.2a shows how the TA from Figure 4.1b was transformed into an STA using Property 5. The idea behind this transformation is that any livelock-accepting execution of T will be mapped to an execution of  $T^+$  that is captured by the new state g. The new g state has an impact on the size of the product (Figure 4.2b), but the strongly connected components of this new product no longer mix non-stuttering transitions and livelock-accepting cycles: this renders the second-pass of Algorithm 3 useless. However, adding an artificial state in STA will increase the size of the product automaton to be explored by the emptiness check (i.e. the first-pass of Algorithm 2). In the next chapter, we present a new improvement of TA, called TGTA that removes the second-pass without adding an artificial state.

The objective of STA is to isolate in the product the exploration of the parts that are composed only by livelock-accepting states and stuttering transitions, like the bold part of the product represented in the Figure 4.2b. In this kind of sub-products, it is easy to find all livelock accepting cycles.

The STA emptiness check algorithm is the first-pass of the TA emptiness check algorithm without the second-pass procedure. In other words, in STA approach, the emptiness check is only Algorithm. 2 (page 49) without line 6.

#### 4.3.3 Correctness of the one-pass emptiness check using STA

In the following lemmas,  $\mathcal{K}, \mathcal{T}, \mathcal{T}^+$  denote respectively a Kripke structure, a TA and an STA.

The first-pass is an SCC-based algorithm, it computes the set of all MSCCs (i.e., Maximal SCCs) of the product automaton. Therefore, in order to prove that the first-pass is sufficient to detect all livelock-accepting cycles, we prove that in  $\mathcal{K} \otimes \mathcal{T}^+$ , searching for all livelock-accepting cycles is equivalent to searching for all MSCCs that are only composed of stuttering transitions and livelock-accepting states. In Algorithm. 2, line 17 allows to detect this kind of MSCCs.

**Lemma 1.** In a product automaton  $\mathcal{K} \otimes \mathcal{T}$ : if one MSCC M contains a product state (s,q) such that q is a livelock-accepting state that has no successors in  $\mathcal{T}$ , then M is only composed of stuttering transitions and livelock-accepting states.

*Proof.* q has no successors in the TA  $\mathcal{T}$ , therefore from q, a run of  $\mathcal{T}$  can only execute stuttering transitions: it stays in the same livelock-accepting state q. Consequently, all product states of M are connected by stuttering transitions. In addition, they have the same livelock-accepting state as TA component (q), therefore by Definition 25 all states of M are livelock-accepting.

**Lemma 2.** In a product automaton  $\mathcal{K} \otimes \mathcal{T}^+$ : one MSCC *M* contains a livelock-accepting state *if and only if M is only composed of stuttering transitions and livelock-accepting states.* 

*Proof.* ( $\Longrightarrow$ ) If an MSCC *M* contains a livelock-accepting state (s,q) of  $\mathcal{K} \otimes \mathcal{T}^+$ , then *q* is a livelock-accepting state that has no successors in  $\mathcal{T}^+$  because in an STA every livelock-accepting state has no successors. The proof follows from Lemma 1 applied to  $\mathcal{K} \otimes \mathcal{T}^+$ .

 $(\Leftarrow)$  Any state of *M* is livelock-accepting.

The difference between lemma 1 and lemma 2 is that the livelock-accepting states of STA have no successors, while those of a TA can. Therefore, the following lemma is true only for STA.

**Lemma 3.** In the product automaton  $\mathcal{K} \otimes \mathcal{T}^+$ : there exists at least one livelock-acceptance cycle *C* if and only if there exists at least one non trivial MSCC *M* such that *M* is only composed of stuttering transitions and livelock-accepting states.

*Proof.*  $(\Longrightarrow)$ : *C* contains at least one livelock-accepting state, therefore applying Lemma 2 with *M* is the MSCC containing *C* allows us to conclude.

( $\Leftarrow$ ): *M* is non-trivial (it contains at least two states or a single state with a self-loop), therefore *M* contains at least one non-trivial cycle only composed of stuttering transitions and livelock-accepting states. This cycle is the livelock-accepting cycle *C*.

In Algorithm. 2, the first-pass computes all MSCCs and line 17 allows to detect only the MSCCs satisfying Lemma 3. Therefore the STA emptiness check algorithm reports one cycle if and only if this cycle is a livelock-accepting cycle or a Büchi-accepting cycle.

**Emptiness check optimizations** According to Lemma 2, it is sufficient to verify that an MSCC contains a livelock-accepting state at line 17 of Algorithm. 2. Therefore, computing SCC.top().k is not necessary for detecting livelock-accepting cycles.

#### 4.3.4 STA optimization

The goal of this optimization is to reduce the number of transitions in STA, by exploiting a special property of livelock-accepting states that are also Büchi-accepting. We begin by introducing a definition to distinguish these particular states:

**Definition 27** (Fully-accepting state). In a TA (or in a product of a TA with a Kripke structure), a state that is both livelock-accepting and Büchi-accepting, is called fully-accepting state.

Definition 26 previously proposed for STA introduces a constraint on livelock-accepting states. This constraint (which we call in the following "livelock-constraint") allowed us to remove the second-pass of the emptiness check. However, in order to transform a TA into an STA that satisfies this "livelock-constraint", we have to add an artificial state g and artificial transitions having g as destination (Property 5).

In this optimization, we start by introducing Property 6 that allows us to deduce that the fullyaccepting states do not require the second pass of the emptiness check. Then, we propose a new STA definition that removes the "livelock-constraint" for the fully-accepting states and therefore reduces the number of artificial transitions added during the transformation of a TA into an STA.

The goal of the following Lemma 4 and Property 6 is is to prove that every cycle C that contains a fully-accepting state  $s_f$  is an accepting cycle. In addition, C is necessarily included in an MSCC M detected by the first-pass of the emptiness check.

We remind that in a product of a TA with a Kripke structure, the definition of an accepting cycle does not depend only on its states but also on its transitions:

• On the one hand, any cycle that contains a livelock-accepting state, is considered as an accepting cycle if it is only composed of stuttering transitions.

• On the other hand, a cycle that contains a Büchi-accepting state, must contains a nonstuttering transition, to be considered an accepting cycle.

In the following lemma 4, we distinguish a particular kind of cycles that can be reported as accepting without checking if they contain or not non-stuttering transitions:

**Lemma 4.** In a product automaton  $\mathcal{K} \otimes \mathcal{T}$ : *if* a cycle *C* contains a fully-accepting state  $s_f$ , *then C* is an accepting cycle.

*Proof.* There are two cases depending on the transitions composing the cycle C:

- If C contains at least one non-stuttering transition, then using the fact that  $s_f$  is Büchi-accepting, we deduce that C is a Büchi-accepting cycle.
- Otherwise, if C contains only stuttering transitions, in this case we exploit the fact that  $s_f$  is a livelock-accepting state to deduce that C is a livelock-accepting cycle.

In addition, these particular accepting cycles also verify the following property:

**Property 6.** In a product automaton  $\mathcal{K} \otimes \mathcal{T}$ : for every accepting cycle *C* containing a fullyaccepting state  $s_f$ , the MSCC *M* that includes *C* is detected by the first-pass of Algorithm. 2.

*Proof.* There are two cases depending on the transitions of the MSCC *M*:

- If *M* includes *C* and contains at least one non-stuttering transition, then *M* is detected by the line 16 of the first-pass, because *M* contains a Büchi-accepting state  $s_f$  and *M* contains non-stuttering transitions.
- Otherwise, if M includes C and is only composed of stuttering transitions, then M is detected by the line 17 of the first-pass, because M contains a livelock-accepting state  $s_f$  and M is only composed of stuttering transitions).

We deduce from Property 6 that the "livelock-constraint" is not necessary for fully-accepting states. This leads us to propose the following optimized definition of STA:

**Definition 28.** A Single-pass Testing Automaton (STA) is a TA  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F}, \mathcal{G} \rangle$  over  $\Sigma$  such that  $\delta \cap ((\mathcal{G} \setminus \mathcal{F}) \times 2^{\Sigma} \times Q) = \emptyset$ . In other words, an STA is a TA in which every state in  $\mathcal{G} \setminus \mathcal{F}$  has no successors.

Consequently, during the TA to STA transformation described by Property 5, it was unnecessary to add artificial transitions (q,k,g) for any transition (q,k,q') where  $q' \in (\mathcal{G} \cap \mathcal{F})$  (i.e., q' is a fully-accepting state).

This optimization of Property 5 is formalized by the following Property 7:

**Property 7.** Let  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F}, \mathcal{G} \rangle$  be a TA, the equivalent STA is  $\mathcal{T}' = \langle Q', I', U', \delta_{\oplus}', \mathcal{F}, \mathcal{G}' \rangle$  where

- $\mathcal{G}' = (\mathcal{G} \cap \mathcal{F}) \cup \{g\}$  where  $g \notin Q$  is a new state,
- $Q' = (Q \setminus G_0) \cup \{g\}$  where  $G_0 = \{q \in G \mid (\{q\} \times \Sigma \times Q) \cap \delta = 0\}$  is the set of states of G that have no successors,



Figure 4.3: Transformation of a TA recognizing a UGb into an optimized STA (according to Property 7). The state 4 is a fully-accepting state

- $I' = I \cup \{g\}$  if  $\mathcal{G} \cap I \neq \emptyset$ , I' = I otherwise,
- $\delta'_{\oplus} = (\delta \setminus (Q \times \Sigma \times G_{\emptyset})) \cup \{(q,k,g) \mid (q,k,q') \in \delta, q' \in ((G \setminus F) \cup G_{\emptyset})\},\$
- $\forall q \in I, U'(q) = U(q) \text{ and } U'(g) = \bigcup_{q \in (\mathcal{G} \cap I)} U(q).$

is such that  $\mathscr{L}(\mathcal{T}') = \mathscr{L}(\mathcal{T}).$ 

Property 7 is illustrated by figure 4.3 that shows the transformation of a TA into an STA satisfying Definition 28. The state 4 is a fully-accepting state in the TA, therefore in the STA we don't add an artificial transition from state 4 to g.

# 4.4 Experimental evaluation of the TA improved emptiness check and of STA

This section presents our experimentation conducted under the same conditions as the previous chapter (see section 3.6), i.e., within the same tools Spot and CheckPN and using the same benchmark inputs (formulas and models).

#### 4.4.1 Implementation

Figure 4.4 shows the building blocks we used to implement the different approaches. The TGBA and BA approaches share the same synchronous product and emptiness check, while a dedicated algorithms are required by the TA and STA approaches.

For TA approach, we have implemented the first-pass improvements of section 4.2 that allow to avoid the second-pass in more cases than the implementation of the previous chapter.

For STA approach, our STA emptiness check implementation shares the same first-pass with the TA algorithm and then disables the second-pass. We also disable the heuristic [46] livelock detection for the STA because it is useless.



Figure 4.4: The experiment's architecture. The transformation of TA into STA is represented by the block "TA2TA<sup>+</sup>". Two command-line switches control which one of the different approaches is used to verify an LTL formula on a Kripke structure.

#### 4.4.2 Results

Table 4.1 and Table 4.2 shows how the TGBA, TA and STA approaches deal with the toy models and random formulas already presented in Section 3.6.2. We omit data for BA since they are outperformed by TGBA (see Section 3.6.3). Table 4.5 and Table 4.6 show toy models against weak-fairness formulas.

Table 4.3 and Table 4.4 show the results of the two cases studies against random, weak-fairness, and dedicated formulas issued from the studies.

All values shown in tables are averaged over 100 different formulas: we checked each model against 100 random and 100 weak-fairness formulas that had no counterexample, and against 100 random and 100 weak-fairness formulas that had a counterexample. The average and maximum are computed separately on each model against each set of formulas.

The cases with formulas verified are separated from cases with violated ones. For verified formulas (Table 4.1, 4.5 and 4.3), no counterexample is found and the full state-space has to be explored. For violated formulas (Table 4.2, 4.6 and 4.4) the on-the-fly exploration of the state-space stopped as soon as the existence of a counterexample could be computed.

Column-wise, these tables show the average and maximum sizes (states and transitions) of: (1) the automata  $A_{\neg \varphi_i}$  expressing the properties  $\varphi_i$ ; (2) the products  $A_{\neg \varphi_i} \otimes \mathcal{K}$  of the property with the model; and (3) the subset of this product that was actually explored by the emptiness check. For verified properties, the emptiness check of TGBA and STA always explores the full product, so the sizes shown in the columns of the product and the emptiness check are equal, while the emptiness check of TA must in many cases perform two passes on the full product (see the linear cloud of green crosses below the diagonal of Figure 4.5a), so the columns of the product and the emptiness check in TA approach show different values. On violated properties, the emptiness check aborts as soon as it finds a counterexample, so the explored size is usually significantly smaller than the full product.

The emptiness check values show a third column labeled "T": this is the time in  $\frac{1}{100^e}$  of seconds (a.k.a. centiseconds) spent doing that emptiness check, including the on-the-fly computation of the subset of the product that is explored. The time spent constructing the property automata from the formulas, shown in column "T<sub> $\phi$ </sub>" (in centiseconds), is negligible compared to that of the emptiness check.

			Verified properties (no counterexample)									
			1	Automator	n	Full	product	Ei	nptiness check			
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т		
	TGBA	avg	12	460	1	2753024	14166594	2753024	14166594	18656		
<b>D5</b>	TODA	max	54	2816	6	11010896	108 553 439	11010896	108553439	80455		
LSO	ТА	avg	83	4324	12	2600539	10293714	4498653	17832896	30334		
ete	IA	max	509	37096	240	9384663	42614845	18769326	85229690	122867		
	STA	avg	83	5131	13	2620919	10388576	2620919	10388576	17784		
	SIA	max	509	48968	292	9867135	45471783	9867135	45471783	65036		
	TCDA	avg	12	411	0	1 3 9 8 0 2 1	10139043	1 398 021	10139043	2325		
	TGBA	max	49	3788	4	4269384	50547840	4269384	50547840	10181		
l go	ТА	avg	69	3118	9	1097667	6606130	2155839	12981954	3744		
Ri	IA	max	509	25773	150	2851152	19625864	5702304	39251728	11149		
	STA	avg	70	3768	10	1229116	7453297	1229116	7453297	2238		
	SIA	max	509	33578	170	2950460	20815264	2950460	20815264	5635		
	тера	avg	9	245	0	2038054	15037218	2038054	15037218	3017		
	IGDA	max	28	2612	4	9132417	89397363	9132417	89397363	15594		
IS5	T۸	avg	48	1643	5	1401286	11000449	1496882	11732337	3084		
E	IA	max	282	11007	98	6109887	54324903	6368802	54324903	14894		
	CTA	avg	49	1962	6	1413836	11090638	1413836	11090638	2889		
	517	max	283	13230	118	6109887	54324903	6109887	54324903	12928		
	TGBA	avg	8	166	0	3356053	32800737	3356053	32800737	5253		
N		max	60	1994	4	20253072	258315134	20253072	258315134	44 28 1		
bar	ТЛ	avg	44	1659	3	2433853	23389805	2433853	23389805	4871		
an	IA	max	277	21251	80	15272712	161 364 553	15272712	161364553	33129		
1	STA	avg	44	1873	4	2433853	23389805	2433853	23389805	4824		
	5171	max	277	22324	87	15272712	161 364 553	15272712	161 364 553	31731		
	TGRA	avg	14	636	1	4668178	27564474	4668178	27564474	9352		
	TODA	max	81	5397	6	17947837	119545256	17947837	119545256	40273		
101	ТА	avg	71	4092	20	2334154	19893200	2334154	19893200	7742		
Phi	IA	max	412	55321	235	8378151	74240975	8378151	74240975	27233		
.	STA	avg	72	4904	22	2334154	19893200	2334154	19893200	7593		
	SIA	max	413	67836	260	8378151	74240975	8378151	74240975	27799		
	TGBA	avg	15	664	0	2682881	30981263	2682881	30981263	9192		
<u>v</u>	TODA	max	82	5792	4	13129644	225905117	13129644	225905117	62328		
lii	ТΔ	avg	95	6048	15	2027251	17168855	3928111	33 295 175	12207		
Rot	1/1	max	481	43 547	149	8169472	68645888	16338944	137291776	48584		
	STA	avg	96	7158	17	2303685	19603937	2303685	19603937	7257		
	5111	max	482	47171	175	8755200	75356160	8755200	75356160	27955		

Table 4.1: Comparison of the three approaches on toy examples with **random formulæ**, when **counterexamples do not exist.** 

			Violated properties (a counterexample exists)										
				Automator	1	Full	product	Ei	nptiness check				
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т			
	TGRA	avg	14	646	0	8118953	42865329	600 828	2325266	3892			
D5	TODA	max	47	3 3 9 2	5	50297469	364 396 821	3632979	22095045	24665			
[ISO]	TA	avg	90	5683	15	8592222	33972863	546562	1886029	3683			
ete		max	656	62947	134	56166663	224 108 070	3178506	11478151	22221			
	STA	avg	91	6824	17	9035627	35735598	590613	2053508	3955			
		max	657	69867	146	56483288	225386395	3644229	13374186	25 569			
	тсра	avg	17	713	1	3163763	26260802	854229	5805080	1 3 2 9			
	TODA	max	51	3788	4	9855589	108310156	5914415	65020741	13999			
ng6	ТА	avg	125	7462	15	2846468	18131253	646 844	3858721	1198			
Ri		max	472	41068	162	8960960	62840996	4923301	36659892	10659			
	STA	avg	126	8960	17	3160137	20284936	728890	4365527	1291			
	51A	max	473	52838	180	9204480	67482964	5041865	38542138	10459			
	тсра	avg	14	496	1	10174866	95061682	1602719	10636471	2148			
	TODA	max	96	4200	3	40672843	404 072 123	11075871	110774947	18644			
1S5	ТА	avg	90	4112	9	8418689	70784299	1044896	7 693 244	1968			
E		max	390	22702	78	31864749	267656316	8518894	63178516	16437			
	STA	avg	91	4855	11	9021110	75667584	1033999	7665018	1979			
	SIA	max	391	26803	84	34082475	286582247	7336271	61 123 648	15463			
	TGBA	avg	13	355	0	10952627	119272618	1153850	8728297	1473			
N		max	87	2184	3	43 0 38 0 8 3	587 543 197	13023468	185410581	26750			
bai	ТА	avg	75	2870	6	8 4 2 9 9 4 8	83120910	670172	5871206	1260			
Kan		max	473	30375	94	31602068	331 833 407	7673335	77 363 942	15562			
	STA	avg	76	3353	7	9248939	90946642	670376	5873036	1241			
	5111	max	474	38397	99	34070884	364831218	7673335	77 363 941	14607			
	TCRA	avg	13	590	0	19118012	195258228	1518296	7 522 623	2651			
		max	74	5928	3	84068722	1 377 479 362	9032250	53881112	15880			
101	ТА	avg	92	5690	14	16012398	148019289	900960	6499587	2558			
Phi		max	356	46498	164	66535322	696331784	7342016	64470840	23 509			
	STA	avg	93	6676	16	17277606	159696150	870692	6437341	2528			
	SIA	max	357	55042	187	70986407	759461806	7 3 4 0 8 0 1	64469273	23022			
	TGBA	avg	21	965	1	5514573	67871153	1676314	18343280	5475			
N		max	74	5928	7	15127431	281 880 971	13627374	191 356 248	52962			
l'ii	ТА	avg	164	11366	21	4869334	41 343 229	1248873	10418682	4110			
Rot		max	854	62419	115	11791872	104488704	7767216	74241202	27244			
	STA	avg	165	13727	23	5452751	46556753	1497531	12536773	4671			
	SIA	max	855	80513	134	12072448	108812032	8860848	85 541 043	30942			

Table 4.2: Comparison of the three approaches on toy examples with **random formulæ**, when **counterexamples exist**.

				Verified properties (no counterexample)										
				1	Automator	1	Full	product	Er	nptiness check				
				st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т			
		TCBA	avg	16	715	1	109158	328976	109158	328976	601			
		TUDA	max	136	5792	4	373 039	2138028	373 039	2138028	2207			
	Ę	Т۸	avg	94	5768	17	94 283	$\mathbf{205464}$	158401	346203	867			
	R	IA	max	481	43 5 47	183	261318	604061	522636	1 208 1 22	2753			
		STA	avg	95	6854	20	103955	227 203	103955	227203	<b>570</b>			
<b>JyORB 3/3/2</b>		SIA	max	482	47171	248	279251	687246	279251	687246	1534			
		тера	avg	5	72	0	96435	277426	96435	277426	<b>524</b>			
		TUDA	max	18	337	2	368 898	1450005	368 898	1450005	2049			
	air	ΤA	avg	79	1106	1	120479	262456	121096	263776	652			
Po	M	IA	max	351	5920	8	633 208	1408943	633 208	1408943	3359			
		CT A	avg	80	1145	1	120483	262464	120483	262464	651			
		SIA	max	352	6072	8	633 208	1408943	633 208	1408943	3373			
		TGBA	. –	7	576	1	345 241	760491	345241	760491	1688			
	$\Phi_{\rm l}$	TA	-	80	14590	8	342613	742815	685226	1485630	3374			
		STA	-	81	17110	12	348 499	760788	348499	760788	1716			
		TGBA	avg	10	333	1	2808973	25212509	2808973	25212509	5325			
			max	64	3813	5	13636352	147 555 158	13636352	147 555 158	28979			
	Ģ	ΤA	avg	50	2239	9	1724291	20203618	1724291	20203618	4852			
	R	IA	max	336	22614	149	8469258	108847708	8469258	108847708	25029			
		CT A	avg	51	2702	10	1724291	20203618	1724291	20203618	4813			
		SIA	max	337	26522	170	8469258	108847708	8469258	108847708	25445			
<b>K</b> 8		тсра	avg	5	39	0	3898645	34822961	3898645	34822961	7344			
<b>P</b>		TUDA	max	21	198	2	14452198	162156912	14452198	162156912	31117			
Z	air	T۸	avg	44	407	1	2112810	24492299	2112810	24492299	6126			
	M	IA	max	171	1920	3	6110748	75624744	6110748	75624744	20653			
		STA	avg	44	422	1	2112810	24492299	2112810	24492299	6067			
		SIA	max	172	2012	3	6110748	75624744	6110748	75624744	18566			
		TGBA	_	6	165	1	46494	302350	46494	302350	50			
	$\Phi_2$	TA	-	9	293	2	33 376	289235	33 376	289235	50			
		STA	-	10	415	1	33 376	289235	33 376	289235	51			

Table 4.3: Comparison of the three approaches for the **case studies when counterexamples do not exist.** 

			[			I	/iolated prope	erties (a countere	example exis	ts)	
				L	Automator	1	Full	product	Eı	nptiness check	
				st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т
		тера	avg	14	760	1	132232	380305	53901	128960	291
		IUDA	max	65	7004	9	1295661	5045814	192692	514905	1049
	P	Т۸	avg	81	6223	29	135237	296081	56518	126300	311
	R	IA	max	540	59971	574	1 377 784	2976709	243950	547459	1284
		CT A	avg	82	7771	33	148699	325993	58989	131 829	324
3/2		SIA	max	541	72599	759	1438600	3109093	253761	570216	1373
3 3/		TGBA	avg	5	66	0	96 849	258704	69053	177391	372
R	WFair		max	22	516	2	432240	1584372	193264	1 0 3 1 8 9 8	1110
<b>JyC</b>		ТА	avg	82	1157	1	201 879	434747	113705	250573	615
$\mathbf{P}_{0}$		IA	max	262	4636	5	560685	1 194 198	254328	594751	1358
		STA	avg	83	1 197	1	202 820	436769	114335	251969	617
		5111	max	263	4674	5	560685	1 194 198	254328	594751	1 392
		TGBA	avg	13	513	1	27266223	405 363 157	1850695	12575125	2702
			max	59	4298	2	91059214	1963331216	13284206	111916150	21771
	P	Т۸	avg	88	4818	13	22458466	298950743	965567	9245311	2392
	$\mathbb{R}$	IA	max	334	35401	165	90049281	1 300 904 178	5241327	61655512	14227
		STA	avg	89	5658	15	23784395	315527818	967 140	9251558	2389
		SIA	max	335	42297	184	91186725	1 3 2 8 3 3 3 5 3 4	5241327	61655512	14256
<b>K</b> 8		TCRA	avg	6	81	0	20904101	300416331	1962765	13896236	3129
AP		TODA	max	23	414	2	101813441	1681905871	9463562	72266873	16928
Z	Fair	Т۸	avg	88	1212	1	27535941	358622897	1112324	11462361	3002
	M	IA	max	245	3749	5	93441682	1258076980	5203746	59642346	16992
		ST4	avg	89	1252	1	28387255	372 261 767	1110154	11456800	3007
		SIA	max	246	3796	4	94454379	1 279 166 220	5 203 741	59642337	15180

Table 4.4: Comparison of the three approaches for the case studies when counterexamples exist.

			Verified properties (no counterexample)									
			A	Automator	ı	Full	product	Ei	mptiness check			
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т		
	TGBA	avg	4	52	0	3071854	15745483	3071854	15745483	21163		
2	TODA	max	17	357	2	9551018	87453872	9551018	87453872	70107		
	ТА	avg	50	564	1	3018533	11771410	3436223	13401691	23 597		
ete	IA	max	184	3362	4	10750205	42228050	11867218	47116180	78154		
	STA	avg	51	589	1	3018776	11772694	3018776	11772694	20870		
	SIA	max	185	3456	4	10750205	42228050	10750205	42228050	77007		
	TGBA	avg	5	64	0	1 1 2 1 5 9 0	7920270	1121590	7920270	1903		
	TODA	max	25	357	2	3873576	38047692	3873576	38047692	7910		
ng(	ТΔ	avg	58	682	1	1118511	6494556	1774991	10328239	3046		
Ri	173	max	183	3353	3	3 3 8 3 3 7 2	22481284	6766744	44962568	11037		
	STA	avg	59	712	1	1126059	6547031	1126059	6547031	2024		
	SIA	max	184	3447	4	3386972	22595612	3386972	22595612	6062		
	TCPA	avg	4	41	0	3030750	23 589 355	3030750	23 589 355	4674		
	TODA	max	12	212	2	11882973	114787553	11882973	114787553	23008		
IS5	ТА	avg	48	515	1	2259230	17600147	2259230	17600147	4589		
E	IA	max	171	4256	3	9764223	79707273	9764223	79707273	18521		
	STA	avg	49	534	1	2259230	17600147	2259230	17600147	4561		
	SIA	max	172	4367	4	9764223	79707273	9764223	79707273	18365		
	TGBA	avg	4	33	1	3098615	29435137	3098615	29435137	4712		
l v		max	21	168	2	13014212	187004177	13014212	187004177	26143		
bal	ТА	avg	36	306	1	2083777	19156589	2083777	19156589	4112		
(an	іл	max	140	2049	3	11039112	106534190	11039112	106534190	23453		
	STA	avg	36	315	1	2083777	19156589	2083777	19156589	4059		
	5171	max	141	2120	3	11039112	106534190	11039112	106534190	22420		
	TGRA	avg	4	35	0	4255979	26793678	4255979	26793678	9064		
0	TODA	max	21	216	2	12974557	112 561 242	12974557	112561242	39806		
101	ТΔ	avg	47	501	1	3229690	27631036	3 2 2 9 6 9 0	27631036	10439		
Phi	173	max	289	8199	6	10987384	94141317	10987384	94141317	38378		
	STA	avg	48	524	1	3229690	27631036	3 2 2 9 6 9 0	27631036	10576		
	5171	max	290	8346	6	10987384	94141317	10987384	94141317	34352		
	TGBA	avg	5	59	0	1782133	17996764	1782133	17996764	5749		
<u>n</u>	TODA	max	16	337	1	5029881	70344648	5029881	70344648	19263		
ii	ΤΔ	avg	54	618	1	1508082	12171559	1802576	14591484	5731		
Rot		max	239	3706	4	3944448	32391168	7 098 368	61216768	22074		
	STA	avg	54	642	1	1520057	12265094	1520057	12265094	4897		
	511	max	240	3763	4	3944448	32391168	3944448	32391168	12200		

Table 4.5: Comparison of the three approaches on toy examples with **weak-fairness formulæ**, **when counterexamples do not exist.** 

					I	/iolated prope	rties (a counter	example exist	s)	
			1	Automator	ı	Full p	product	En	nptiness check	
			st.	tr.	$T_{\boldsymbol{\phi}}$	st.	tr.	st.	tr.	Т
	TGBA	avg	5	78	0	3731947	17751951	522414	1759041	3438
n5	TODA	max	18	351	2	13508730	79589415	3696481	16733420	24769
LSO	ТА	avg	95	1400	2	10950499	42790086	641997	2123401	4370
ete		max	350	5968	13	30354888	120985972	7637347	29621420	51473
P	STA	avg	95	1448	1	10951464	42793386	642215	2124081	4362
	SIA	max	351	6120	9	30354888	120985972	7637348	29621421	51487
	TGBA	avg	7	123	1	2000787	15868857	890 567	5736568	1 410
5	TODA	max	38	1168	3	4670124	53246768	2766283	26917302	6101
ng(	ТА	avg	128	2009	2	3456785	21069783	905265	5156327	1621
Ri		max	388	12783	14	6152232	43283696	2747578	17464827	4793
	STA	avg	129	2063	2	3459344	21101063	907 368	5174029	1614
	517	max	389	12873	14	6160248	43438480	2747578	17464827	5352
	TCPA	avg	5	68	0	7978129	75701408	1372877	7 544 397	1607
	TUDA	max	23	370	2	27 570 312	293219919	6676584	53451994	9961
<b>1</b> S5	TA	avg	71	871	1	10007988	81736556	848557	5707348	1522
FN		max	245	3766	5	27733464	249573778	6425987	51962113	12307
	STA	avg	71	900	1	10128248	82964077	841164	5667421	1511
	SIA	max	246	3766	4	27733464	249 573 778	6425987	51962113	11806
	TGBA	avg	5	66	0	7040236	73867374	1 344 870	9279129	1562
2		max	17	305	2	25423161	388961954	7950633	80041971	11336
bar	тл	avg	69	849	1	8102548	77515606	809467	6987095	1474
an		max	237	3655	3	32732392	329 125 090	5 501 242	53375472	10218
×	STA	avg	70	875	1	8355726	80541627	788150	6773391	1442
	SIA	max	238	3710	3	33147912	335772570	5498955	53 371 193	10147
	TGRA	avg	5	69	0	10515366	102622810	833663	3762217	1357
0	IUDA	max	25	392	3	24496745	402197001	5017174	32273634	9802
101	ТА	avg	81	1027	1	16779360	150030607	742819	5019286	1943
Phi		max	267	4339	5	43186049	394172486	6703064	55546036	20284
	STA	avg	81	1068	1	16801072	150313327	727226	4926991	1920
	SIA	max	268	4386	6	43 194 029	394299774	6703064	55546036	20201
	TGBA	avg	9	157	1	4006015	46910526	1486839	14360747	4691
N		max	44	1168	2	14802883	195152925	5666009	78112291	21619
in	ТА	avg	143	2467	2	5832674	48201707	1348406	10788993	4327
Rot		max	466	12783	14	18204684	151 155 748	6486142	53115008	18844
	STA	avg	144	2529	2	5841364	48285515	1354923	10843123	4324
	SIA	max	467	12873	13	18204684	151 155 748	6486142	53115008	19093

Table 4.6: Comparison of the three approaches on toy examples with **weak-fairness formulæ**, **when counterexamples exist.** 



(b) STA against TGBA approaches

Figure 4.5: Performance (number of transitions explored by the emptiness check) of STA against TA and TGBA.

Figure 4.5 compares the number of visited transitions when running the emptiness check; plotting STA against TA and TGBA. This gives an idea of their relative performance. Each point corresponds to one of the 5600 evaluated formulas (2800 violated with counterexample as black circles, and 2800 verified having no counterexample as green crosses). Each point below the diagonal is in favor of STA while others are in favor of the other approach. Axes are displayed using a logarithmic scale.

No comparison is presented with BA since they are less efficient than TGBA (see the experimental evaluation presented in section 3.6.3).

All these experiments were run on a 64bit Linux system running on an Intel(R) 64-bit Xeon(R) @2.00GHz, with 10GB of RAM.

#### 4.4.3 Discussion

Before discussing the performance of different approaches, we recall that in our implementation using CheckPN tool, the cost of computing labels in the Kripke structure is higher (despite the fact that we use a cache). This increases the computing time of products in TA and STA approaches (more than in the approaches TGBA and BA), because TA and STA approaches query two labels by transition of the Kripke structure (to compute an xor between source label and destination label) while other approaches query only one label.

In an implementation where computing labels is cheap, the execution time should be proportional to the number of transitions explored by the emptiness check, so it is important to not consider only the execution time provided by our experiments.

**On verified properties**, the results are very straightforward to interpret when looking at the number of states and transitions explored by the emptiness check.

The emptiness check optimizations proposed in Section 4.2 improve the performance of the TA approach. This can be observed by comparing the results in this chapter and the results of TA approach in the previous chapter (see Section 3.6.3: for verified properties, the number of transitions visited by the emptiness check was always twice the number of transitions in the product).

TA outperform TGBA except for both Random and weak-fairness properties against Peterson, Ring, Robin and PolyORB.

STA significantly improve TA in all cases where a second pass was necessary. In these cases, the STA approach, with its single-pass emptiness check, is a clear improvement over TA. These cases where the STA approach is twice faster than TA's, appear as a linear cloud of green crosses below the diagonal in the scatter plot of Figure 4.5a (we recall that the axes are displayed using a logarithmic scale) Otherwise, they have the same performance because if no livelock-acceptance states are detected in the product then the TA and STA approaches explore exactly the same product (these cases correspond to the green crosses on the diagonal).

In the scatter plot comparing STA against TGBA, the green crosses appear on both side of the diagonal, with much more points where STA is better. Furthermore, in the results tables, if we observe in more details the average number of states/transitions explored during the emptiness check, STA outperform TGBA in all cases except for weak-fairness formulas against Ring and PolyORB. In this cases, TGBA benefits from the large number of acceptance conditions generated when translating weak-fairness formulas.

**On violated properties**, it is harder to interpret the results because they depend on the order in which non-deterministic transitions of the property automaton are explored. In the best case, the order of transitions leads the emptiness check straight to a counterexample; in the worst case, the algorithm explores the whole product until it finally finds a counterexample.

The different kinds of property automata TGBA, TA and STA provide different orders of transitions and therefore change the number of states and transitions to be explored by the emptiness check before a counterexample is found.

If we analyze in more details the results tables, we observe that:

- For the (full) product size, TA and STA produce a smaller product on the average than TGBA for random formulas. However, for weak-fairness formulas, TGBA produces the smallest product on the average.
- For the emptiness check, looking at the average number of transitions explored by the emptiness check (and taking into account the extra computation of labels discussed previously), indicates that TA and STA approaches outperform the TGBA approach except for weak-fairness formulas against Peterson, Philo and PolyORB. Even for STA where an artificial state and non-deterministic transitions are added, no significant overhead is noticed.

## 4.5 Conclusion

In a preliminary work presented in the previous chapter, we experiment LTL model checking of stuttering-insensitive properties with various techniques: Büchi automata (BA), Transitionbased Generalized Büchi Automata and Testing Automata [46]. At this time, conclusions were that TA has good performance for violated properties (i.e. when a counterexample was found). However, this was not the case when no counterexample was computed since the entire product had to be visited twice to check for each acceptance mode of a TA (Büchi acceptance or livelock-acceptance).

This chapter extends the above work in two ways. First, it introduces a modified emptiness check algorithm that avoids the second pass when it is useless (i.e., no livelock-acceptance state is detected in the product). Second, it proposes a transformation of TA into STA that avoids the need for a second pass (in all cases).

Both new algorithms have been implemented in Spot, our model checking library and used on several benchmark models including large models issued from case studies. Experimentation with Spot reported that, STA remain good for violated properties, and also beat TA and TGBA in most cases when properties exhibit no counterexample.

In the next chapter, we introduce a new kind of automata that combines ideas from TGBA and STA. The goal is to have a form of single-pass testing automata without adding any artificial state and with generalized acceptance conditions on transitions like a TGBA. These generalized acceptance conditions allow to improve the performance when checking weak-fairness formulas, as observed in the experimentation with better performance of TGBA for this kind of formulas.

## CHAPTER 5

# Transition-based Generalized Testing Automata (TGTA): A Single-pass and Generalized New Automata

#### Contents

5.1	Introduction										
5.2	Transi	Introduction87Iransition-based Generalized Testing Automata (TGTA)885.2.1Ø-TGTA895.2.2TGTA90IGTA Construction915.3.1From TGBA to Ø-TGTA: Construction of an intermediate Ø-TGTA from a TGBA915.3.2From Ø-TGTA to TGTA: Elimination of useless stuttering-transitions (Ø) without introducing livelock-acceptance93Explicit Model checking using TGTA965.4.1Synchronous Product of a TGTA with a Kripke Structure965.4.2Emptiness check (the same as TGBA)97Experimental evaluation of TGTA98									
	5.2.1	Ø-TGTA	89								
	5.2.2	TGTA	90								
5.3	TGTA	Construction	91								
	5.3.1	From TGBA to Ø-TGTA: Construction of an intermediate Ø-TGTA from a TGBA	91								
	5.3.2	From $\emptyset$ -TGTA to TGTA: Elimination of useless stuttering-transitions ( $\emptyset$ ) without introducing livelock-acceptance	93								
5.4	Explic	it Model checking using TGTA	96								
	5.4.1	Synchronous Product of a TGTA with a Kripke Structure	96								
	5.4.2	Emptiness check (the same as TGBA)	97								
5.5	Exper	imental evaluation of TGTA	98								
	5.5.1	Implementation	98								
	5.5.2	Results	99								
	5.5.3	Discussion	107								
	5.5.4	Experimental Results once the TGBA is improved by simulation-reduction	108								
5.6	Concl	usion	108								

## 5.1 Introduction

In Chapter 3, we have shown that Testing Automata (TA) are better than Büchi Automata when the formula to be verified is violated (i.e., a counterexample is found), but this is not the case when the property is verified since the entire product have to be visited twice to check for each acceptance mode of a TA. Then, in order to improve the TA approach, we proposed STA in Chapter 4. STA is a transformation of TA into a normal form that does not need the second pass of the emptiness

Chapter 5. Transition-based Generalized Testing Automata (TGTA): A Single-pass and 688 Generalized New Automata



Figure 5.1: An overview of the relations between the different variants of  $\omega$ -automata presented in this thesis.

check. Unfortunately, STA can increase the size of the product automaton, because in order to remove the second pass, the transformation from TA into STA adds an artificial state.

This chapter introduces a new type of  $\omega$ -automata for stutter-invariant properties, called Transition-based Generalized Testing Automata (TGTA) [6], that mixes features from both TA and TGBA.

The basic idea of TGTA is to build an improved form of testing automata with generalized acceptance conditions on transitions, which allows us to modify the automata construction in order to remove the second pass of the emptiness check of the product. The constructed TGTA represents all the stuttering-transitions using only self-loops.

TGTA combines the advantages of TA and TGBA and it is better than STA because TGTA allows to remove the second pass without adding an artificial state.

In addition, as seen in Figure 5.1, a TA is built from a BA while a TGTA is built directly from a TGBA. Therefore, compared to TA, TGTA can benefit from the fact that TGBA are smaller than BA. Indeed, TGBA are more concise [49, 36] because they use generalized acceptance conditions on transitions, especially for weak-fairness formulas, as was already observed in Section 3.6.3 page 54.

Another advantage of TGTA compared to TA, is that the implementation of TGTA approach does not require a dedicated emptiness check, it reuses the same algorithm used for TGBA (and BA), and the counterexample constructed by this algorithm is also reported as a counterexample for the TGTA approach. We are thus able to compare TGTA with the "traditional" algorithms we used on TA, BA and TGBA. The results of these experimental comparisons show that TGTA compete well on our examples: the TGTA approach is statistically more efficient than the other evaluated approaches, especially when no counterexample is found (i.e., the property is verified) because it does not require a second pass.

### 5.2 Transition-based Generalized Testing Automata (TGTA)

The following definition of TGTA combines features from both TGBA and TA. From TGBA, we inherit the use of transition-based generalized acceptance conditions. From TA, we take the idea of labeling transitions with changesets, however we remove the use of livelock-acceptance (because it may require a two-pass emptiness check), and we remove the implicit stuttering (in TGTA,  $\delta$  explicitly represents the stuttering transitions, but we will see in the following that our constructed TGTA satisfies a stuttering-normalization constraint that optimizes the representation

of these stuttering transitions).

The resulting Chimera, accepts only stuttering-insensitive languages like TA, and inherits advantages from both TA and TGBA: it has a simple one-pass emptiness-check procedure (the same as the one for TGBA), and can benefit from reductions based on the stuttering of the properties pretty much like a TA. Livelock acceptance states, which are no longer supported, can be emulated using states with stuttering self-loops labeled by the set of all acceptance conditions  $\mathcal{F}$  (these particular self-loops are called "accepting stuttering self-loops" in the following).



Figure 5.2: The two steps of the construction of a TGTA from a TGBA.

#### 5.2.1 Ø-TGTA

Before defining TGTA, we begin by defining an intermediate form called  $\emptyset$ -TGTA ("emptychangesets" TGTA). This intermediate form is a generalized Testing Automaton that allows to represent any LTL property. Then, we define TGTA as a normal form of  $\emptyset$ -TGTA used to only represent stutter-invariant properties (LTL\X). In Section 5.3 we will show that a TGTA is constructed in two steps as illustrated in Figure 5.2. The first step transforms a TGBA into an  $\emptyset$ -TGTA by labeling transitions with changesets. The second step transforms an  $\emptyset$ -TGTA into a TGTA by normalizing the representation of stuttering transitions. This normalization benefits from the hypothesis that the LTL property is stutter-invariant to remove all stuttering transitions that are not self-loops in TGTA.

In addition to being used in the construction of TGTA, 0-TGTA will also be used to represent the product between a TGTA and a Kripke structure in Section 5.4.1.

**Definition 29** (0-TGTA). An 0-TGTA over the alphabet  $\Sigma$  is a tuple  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F} \rangle$  where:

- *Q* is a finite set of states,
- $I \subseteq Q$  is a set of initial states,
- $U: I \to 2^{\Sigma}$  is a function mapping each initial state to a set of symbols of  $\Sigma$ ,
- $\mathcal{F}$  is a finite set of acceptance conditions,
- $\delta \subseteq Q \times \Sigma \times 2^{\mathcal{F}} \times Q$  is the transition relation, where each element (q,k,F,q') represents a transition from state q to state q' labeled by a changeset k interpreted as a (possibly empty) set of atomic propositions whose values change between q and q', and the set of acceptance conditions  $F \in 2^{\mathcal{F}}$ ,

An infinite word  $\sigma = \ell_0 \ell_1 \ell_2 \ldots \in \Sigma^{\omega}$  is accepted by  $\mathcal{T}$  if there exists an infinite path  $r = (q_0, \ell_0 \oplus \ell_1, F_0, q_1)(q_1, \ell_1 \oplus \ell_2, F_1, q_2)(q_2, \ell_2 \oplus \ell_3, F_2, q_3) \ldots \in \delta^{\omega}$  where:

- $q_0 \in I$  with  $\ell_0 \in U(q_0)$  (the infinite word is recognized by the path),
- $\forall f \in \mathcal{F}, \forall i \in \mathbb{N}, \exists j \ge i, f \in F_j$  (each acceptance condition is visited infinitely often).

Chapter 5. Transition-based Generalized Testing Automata (TGTA): A Single-pass and 90 Generalized New Automata



Figure 5.3: An  $\emptyset$ -TGTA (left) and a TGTA (right) for the LTL property  $\varphi = FGp$ , with acceptance conditions indicated by the black dot  $\bullet$ .

#### *The language accepted by* $\mathcal{T}$ *is the set* $\mathscr{L}(\mathcal{T}) \subseteq \Sigma^{\omega}$ *of infinite words it accepts.*

Figure 5.3a shows an  $\emptyset$ -TGTA recognizing the LTL formula FG *p*. Acceptance sets are represented using dots as in TGBAs. Transitions are labeled by changesets: e.g., the transition  $(0, \{p\}, 1)$  means that the value of *p* changes between states 0 and 1. Initial valuations are shown above initial arrows:  $U(0) = \{\bar{p}\}, U(1) = \{p\}$  and  $U(2) = \{p\}$ . Any infinite path in this example is accepted if it visits infinitely often, the acceptance transition indicated by the black dot  $\bullet$ : i.e., the stuttering self-loop  $(2, \emptyset, \bullet, 2)$ . As an illustration, the infinite word  $\bar{p}; p; p; p; \dots$  is accepted by the run  $(0) \xrightarrow{\{p\}} (1) \xrightarrow{\emptyset} (2) \xrightarrow{\emptyset} (2) \cdots$  Indeed, a run recognizing such an infinite word must start in state 0 (because only  $U(0) = \{\bar{p}\}$ ), then it changes the value of *p*, so it has to move to state 1 because from state 0 only the transition  $(0, \{p\}, 1)$  is labeled by  $\{p\}$ . In the next step, the value of *p* does not change and the run must execute a stuttering transition among  $(1, \emptyset, 1)$  or  $(1, \emptyset, 2)$ . To be accepted, it must eventually move to state 2 (rather than remain in state 1), and finally stay on state 2 by executing infinitely the accepting stuttering self-loop  $(2, \emptyset, \bullet, 2)$ .

#### 5.2.2 TGTA

A TGTA is a normal form of  $\emptyset$ -TGTA satisfying a structural constraint on stuttering transitions, called "stuttering-normalization constraint", its objective is to force the TGTA to represent the stuttering transitions using only stuttering self-loops.

**Definition 30** (TGTA). A TGTA is an  $\emptyset$ -TGTA  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F} \rangle$  such that  $\mathscr{L}(\mathcal{T})$  is stutterinvariant and the transition relation  $\delta$  has to satisfy the following stuttering-normalization constraint:

1. All stuttering transitions are self-loops, and

2. and every state has a stuttering self-loop.

Formally, the stuttering-normalization constraint can be expressed by the following equivalence:  $\forall (q,q') \in Q^2 : (\exists F \in 2^{\mathcal{F}}, (q, \emptyset, F, q') \in \delta) \iff (q = q')$ 

Figure 5.3b shows a TGTA recognizing the LTL formula FG p. We note that this TGTA satisfies the stuttering-normalization constraint. Indeed, we have that the set of stuttering transitions is  $\{(0, \emptyset, \emptyset, 0), (1, \emptyset, \emptyset, 1), (2, \emptyset, \bullet, 2)\}$ , i.e., a stuttering self-loop on each state. In this TGTA, the

infinite word  $\bar{p}; p; p; p; \dots$  is accepted by the run  $(\bigcirc^{\{p\}} (2) \stackrel{\emptyset}{\longrightarrow} (2) \stackrel{\emptyset}{\longrightarrow} (2) \cdots)$  because the value p only changes between the first two steps. Indeed, a run recognizing such an infinite word must start in state 0 (because only  $U(0) = \{\bar{p}\}$ ), then it changes the value of p, so it has to take transitions labeled by  $\{p\}$ , i.e.,  $(0, \{p\}, 1)$  or  $(0, \{p\}, 2)$ . To be accepted, it must move to state 2 (rather than state 1), and finally stay on state 2 by executing infinitely the accepting stuttering self-loop  $(2, \emptyset, \bullet, 2)$ .

In this work, we define and use TGTA only for stutter-invariant LTL properties. Indeed, we need this restriction to build a TGTA that satisfies the stuttering-normalization constraint (in particular, we will show in section 5.3.2 how we exploited this restriction to remove all the stuttering transitions that are not self-loops).

In the next section, we present in detail the formalization of the different steps used to build a TGTA that satisfies Definition 30.

### 5.3 TGTA Construction

Let us now describe how to build a TGTA starting from a TGBA of a stutter-invariant LTL property. The construction is inspired by the one presented in Section 3.4 (page 42) that constructs a TA from a BA.

Similar to TA construction, a TGTA is built in two steps as illustrated in Figure 5.2, the first one builds an intermediate  $\emptyset$ -TGTA from a TGBA. Then, the second step builds the final form of TGTA by removing the useless stuttering transitions of the  $\emptyset$ -TGTA. This simplification of stuttering transitions does not introduce livelock-accepting states in TGTA (this represents a crucial difference between TA and TGTA because livelock-accepting states require a second pass in the emptiness check of the product using TA).

Figure 5.4d shows a TGTA constructed for  $a \cup Gb$  in the same way as we did for Figure 3.3d (page 43). The only accepting runs are those that see  $\bullet$  infinitely often. The reader can verify that all the infinite words taken as example in section 3.4 are still accepted, but not always with the same runs (for instance  $ab; \bar{a}b; \bar{a}b; \bar{a}b; \ldots$  is accepted by the run 2,4,4,4,..., but not by the run 2,3,3,3,...). This difference is due to the way we emulate livelock-accepting states, as we will describe later (in Property 9 page 93).

## 5.3.1 From TGBA to 0-TGTA: Construction of an intermediate 0-TGTA from a TGBA

This first step is similar to the first step of the TA construction and the following first property is the counterpart of Property 3 (i.e., transforming a BA into an  $\emptyset$ -TA presented in Section 3.4 page 42). We construct an  $\emptyset$ -TGTA from a TGBA by moving labels to states, and labeling each transition by the set difference between the labels of its source and destination states. While doing so, we keep the generalized acceptance conditions on the transitions. An example of a constructed  $\emptyset$ -TGTA is shown on Figure 5.4b.

**Property 8** (Converting TGBA into 0-TGTA). For any TGBA  $\mathcal{G} = \langle Q_{\mathcal{G}}, I_{\mathcal{G}}, \delta_{\mathcal{G}}, \mathcal{F} \rangle$  over the alphabet  $\Sigma = 2^{AP}$  and such that  $\mathscr{L}(\mathcal{G})$  is stutter-invariant, let us define the 0-TGTA  $\mathcal{T} =$ 



Figure 5.4: TGTA obtained after various steps while translating the TGBA representing  $a \cup Gb$ , into a TGTA with  $\mathcal{F} = \{\bullet\}$ .

 $\langle Q_{\mathcal{T}}, I_{\mathcal{T}}, U_{\mathcal{T}}, \delta_{\mathcal{T}}, \mathcal{F} \rangle \text{ with } Q_{\mathcal{T}} = Q_{\mathcal{G}} \times \Sigma, \ I_{\mathcal{T}} = I_{\mathcal{G}} \times \Sigma \text{ and}$   $(i) \ \forall (q, \ell) \in I_{\mathcal{T}}, U_{\mathcal{T}}((q, \ell)) = \{\ell\}$ 

(*ii*)  $\forall (q,\ell) \in Q_{\mathcal{T}}, \forall (q',\ell') \in Q_{\mathcal{T}}, ((q,\ell),\ell \oplus \ell',F,(q',\ell')) \in \delta_{\mathcal{T}} \iff ((q,\ell,F,q') \in \delta_{\mathcal{G}})$ *Then*  $\mathscr{L}(\mathcal{G}) = \mathscr{L}(\mathcal{T}).$ 

#### Proof of property 8.

( $\subseteq$ ) Let  $\sigma_1 = \ell_0 \ell_1 \ell_2 \ldots \in \mathscr{L}(\mathcal{G})$  be an infinite word accepted by  $\mathcal{G}$ .

By Definition 16,  $\sigma_1$  is recognized by a path  $(q_0, \ell_0, F_0, q_2)(q_2, \ell_1, F_1, q_2) \dots \in \delta_{\mathcal{G}}^{\omega}$  of  $\mathcal{G}$ , such that  $q_0 \in I$ , and  $\forall f \in \mathcal{F}, \forall i \in \mathbb{N}, \exists j \ge i, f \in F_j$ . By applying (ii) and (i), we can see that there exists a corresponding path  $((q_0, \ell_0), \ell_0 \oplus \ell_1, F_0, (q_1, \ell_1))((q_1, \ell_1), \ell_1 \oplus \ell_2, F_1, (q_2, \ell_2)) \dots \in \delta_{\mathcal{T}}^{\omega}$  of  $\mathcal{T}$  such that  $(q_0, \ell_0) \in I_{\mathcal{T}}, \ell_0 \in U_{\mathcal{T}}((q_0, \ell_0))$ , and still  $\forall f \in \mathcal{F}, \forall i \in \mathbb{N}, \exists j \ge i, f \in F_j$ . By Definition 30 we therefore have  $\sigma_1 \in \mathscr{L}(\mathcal{T})$ .

(2) Let  $\sigma_2 = w_0 w_1 w_2 \ldots \in \mathscr{L}(\mathcal{T})$  be an infinite word accepted by  $\mathcal{T}$ .

By Definition 30,  $\sigma_2$  is recognized by a path  $((q_0, \ell_0), w_0 \oplus w_1, F_0, (q_1, \ell_1))((q_1, \ell_1), w_1 \oplus w_2, F_1, (q_2, \ell_2))$ 

 $\dots \in \delta^{\omega}_{\mathcal{T}} \text{ of } \mathcal{T} \text{ such that } (q_0, \ell_0) \in I_{\mathcal{T}}, w_0 \in U_{\mathcal{T}}((q_0, \ell_0)), \text{ and } \forall f \in \mathcal{F}, \forall i \in \mathbb{N}, \exists j \ge i, f \in F_j. \text{ Of course we have } w_i \oplus w_{i+1} = \ell_i \oplus \ell_{i+1} \text{ but this does not suffice to imply that } \ell_i = w_i. \text{ However (i) tells us that } w_0 \in U_{\mathcal{T}}((q_0, \ell_0)) = \{\ell_0\} \text{ so } w_0 = \ell_0, \text{ and since } w_i \oplus w_{i+1} = \ell_i \oplus \ell_{i+1} \text{ it follows that } w_i = \ell_i. \text{ By applying (ii) can now find a corresponding path } (q_0, \ell_0, F_0, q_2)(q_2, \ell_1, F_1, q_2) \dots \in \delta^{\omega}_{\mathcal{G}} \text{ of } \mathcal{G}, \text{ such that } q_0 \in I, \forall i \in \mathbb{N}, (w_i = \ell_i), \text{ and } \forall f \in \mathcal{F}, \forall i \in \mathbb{N}, \exists j \ge i, f \in F_j. \text{ By Definition 16 we therefore have } \sigma_2 \in \mathscr{L}(\mathcal{G}). \qquad \Box$ 

## **5.3.2** From 0-TGTA to TGTA: Elimination of useless stuttering-transitions (0) without introducing livelock-acceptance

The next property is the pendent of Property 4 (page 45) to simplify the automaton by removing stuttering transitions and thus obtain the final form of the TGTA. Here we cannot remove self-loop transitions labeled by  $\emptyset$  (i.e., stuttering self-loops), but we can remove all others. The intuition behind this simplification is illustrated in Figure 5.5a:  $q_0$  is reachable from state q by a non-stuttering transition, but  $q_0$  can reach an accepting stuttering-cycle by following only stuttering transitions. In the context of TA we would have to declare  $q_0$  as being a livelock-accepting state. For TGTA, we replace the accepting stuttering-cycle by adding a self-loop labeled by all acceptance conditions on  $q_n$ , then the predecessors of  $q_0$  are connected to  $q_n$  as in Figure 5.5b.

In the last step of the following construction, in order to maintain the same accepted (stuttering) language, we add a stuttering self-loop to each state before removing all stuttering transitions between every two distinct states.

**Property 9** (Elimination of useless stuttering transitions of  $\emptyset$ -TGTA to build a TGTA). Let  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F} \rangle$  be a  $\emptyset$ -TGTA such that  $\mathscr{L}(\mathcal{T})$  is stutter-invariant. By combining the first three of the following operations, we can remove all stuttering transitions that are not self-loop (see Figure 5.5) and therefore obtain a TGTA. The fourth operation can be performed along the way for further (classical) simplifications.



(a)  $\emptyset$ -TGTA: Before reduction of stuttering transitions ( $\emptyset$ ).



(b) After reduction of stuttering transitions (0).



(c) TGTA: After removing all stuttering transitions and adding stuttering self-loops on all states.

Figure 5.5: Elimination of useless stuttering transitions of Ø-TGTA to obtain a TGTA.

- If Q ⊆ Q is a SCC such that any two states q,q' ∈ Q can be connected using a sequence of stuttering transitions (q, 0, F<sub>0</sub>, r<sub>1</sub>)(r<sub>1</sub>, 0, F<sub>1</sub>, r<sub>2</sub>) ··· (r<sub>n</sub>, 0, F<sub>n</sub>, q') ∈ δ\* with F<sub>0</sub> ∪ F<sub>1</sub> ∪ ··· ∪ F<sub>n</sub> = F, then we can add an accepting stuttering self-loop (q, 0, F, q) on each state q ∈ Q. I.e., the 0-TGTA T' = ⟨Q, I, U, δ ∪ {(q, 0, F, q) | q ∈ Q}, F⟩ is such that L(T') = L(T). Let us call such a component Q an accepting Stuttering-SCC.
- 2. If there exists an accepting Stuttering-SCC Q and a sequence of stuttering-transitions:  $(q_0, \emptyset, F_1, q_1)(q_1, \emptyset, F_2, q_2) \cdots (q_{n-1}, \emptyset, F_n, q_n) \in \delta^*$  such that  $q_n \in Q$  and  $q_0, q_1, \dots, q_{n-1} \notin Q$ (as shown in Figure 5.5a), then:
  - For any non-stuttering transition,  $(q,k,F,q_0) \in \delta$  going to  $q_0$  and such that  $k \neq \emptyset$ (and  $(q,k,F,q_n) \notin \delta$ ), the  $\emptyset$ -TGTA  $\mathcal{T}'' = \langle Q, I,U, \delta \cup \{(q,k,F,q_n)\}, \mathcal{F} \rangle$  is such that  $\mathscr{L}(\mathcal{T}'') = \mathscr{L}(\mathcal{T})$ .
  - If  $q_0 \in I$ , the 0-TGTA  $\mathcal{T}'' = \langle Q, I \cup \{q_n\}, U'', \delta, \mathcal{F} \rangle$  with  $\forall q \neq q_n, U''(q) = U(q)$  and  $U''(q_n) = U(q_n) \cup U(q_0)$ , is such that  $\mathscr{L}(\mathcal{T}'') = \mathscr{L}(\mathcal{T})$ .
- 3. Let T<sup>†</sup> = ⟨Q, I<sup>†</sup>, U<sup>†</sup>, δ<sup>†</sup>, F⟩ be the 0-TGTA obtained after repeating the previous two operations as much as possible (i.e., T<sup>†</sup> contains all the transitions and initial states that can be added by the above two operations (Figure 5.5b)). Then, we can add non-accepting stuttering self-loops (q, 0, 0, q) to all states that did not have an accepting stuttering self-loop (Figure 5.5c), because T describes a stuttering invariant property. Also we can remove all stuttering transitions that are not self-loops since stuttering can be captured by self-loops after the previous two operations. After this last reduction of stuttering transitions, we obtain the final TGTA. More formally, the TGTA T<sup>'''</sup> = ⟨Q, I<sup>†</sup>, U<sup>†</sup>, δ<sup>'''</sup>, F⟩ with δ<sup>'''</sup> = {(q,k,F,q') ∈ δ<sup>†</sup> | k ≠ 0 ∨ (q = q' ∧ F = F)} ∪ {(q,0,0,q) | (q,0,F,q) ∉ δ<sup>†</sup>} is such that L(T<sup>'''</sup>) = L(T<sup>†</sup>) = L(T<sup>†</sup>).

4. Any state from which one cannot reach a Büchi-accepting cycle can be removed from the automaton without changing its language.

Here again, an additional optimization is to merge bisimilar states, this can be achieved using the same algorithm used to simplify a TA, taking Q as initial partition and taking into account the acceptance conditions of the outgoing transitions. All these steps are shown in Figure 5.4.

Proof of property 9.

- (T' ⊇ T) Obvious because we are only adding transitions. (T' ⊆ T) Let δ' = δ ∪ {(q,0, F,q) | q ∈ Q}. Consider an accepting infinite word σ = ℓ<sub>0</sub>ℓ<sub>1</sub>ℓ<sub>2</sub>... ∈ ℒ(T') recognized by an accepting path π' on T'. Any transition of π' that is not in δ is a self-loop (q,0, F,q) that has been added to δ' because an accepting stuttering-SCC exists in δ around q: so any (q,0, F,q) ∈ δ' can be replaced by a sequence of stuttering transitions (q,0,G<sub>0</sub>,r<sub>1</sub>)(r<sub>1</sub>,0,G<sub>1</sub>,r<sub>2</sub>)...(r<sub>n</sub>,0,G<sub>n</sub>,q) ∈ δ\* such that G<sub>0</sub> ∪ G<sub>1</sub> ∪ ...G<sub>n</sub> = F. The path π ∈ δ<sup>∞</sup> obtained by replacing all such transitions is an accepting path of T that recognizes a word that is stuttering equivalent to σ. Since ℒ(T) is stuttering-insensitive, it must also contain σ.
- (𝒯" ⊇ 𝒯) Obvious for the same reason. (𝒯" ⊆ 𝒯) We consider the case where q<sub>0</sub> is non initial (the initial case is similar). Let δ" = δ∪ {(s,k,F,q<sub>n</sub>)}. Consider an accepting infinite word σ = ℓ<sub>0</sub>ℓ<sub>1</sub>ℓ<sub>2</sub>... ∈ ℒ(𝒯") recognized by a path 𝒯" on 𝒯". Let 𝔅 be the path on 𝒯 obtained by replacing in 𝒯" any occurrence of (s,k,F,q<sub>n</sub>) ∈ (δ" \δ) by the sequence (s,k,F,q<sub>0</sub>)(q<sub>0</sub>,Ø,F<sub>1</sub>,q<sub>1</sub>)(q<sub>1</sub>,Ø,F<sub>2</sub>,q<sub>2</sub>)...(q<sub>n-1</sub>,Ø,F<sub>n</sub>,q<sub>n</sub>) ∈ δ\*. The path 𝔅 ∈ δ<sup>ω</sup> is also an accepting path of 𝒯 that recognizes a word that is stuttering equivalent to σ. Since ℒ(𝒯) is stuttering-insensitive, it must also contain σ.
- \$\mathcal{L}(\mathcal{T}^{\dagger}) = \mathcal{L}(\mathcal{T})\$ by application of the previous two properties, therefore \$\mathcal{L}(\mathcal{T}^{\dagger})\$ is a stuttering-insensitive language because \$\mathcal{T}'''\$ is obtained from \$\mathcal{T}^{\dagger}\$ that recognizes a stuttering-insensitive language, by adding stuttering self-loops on all its states before removing all stuttering transitions that are not self-loops.

To prove that two stuttering-insensitive languages are equal, it is sufficient to verify that they contain the same words of the following two forms:

- $\sigma = \ell_0 \ell_1 \ell_2 \dots$  with  $\forall i \in \mathbb{N}, \ell_i \oplus \ell_{i+1} \neq \emptyset$  (non-stuttering words), or
- $\sigma = \ell_0 \ell_1 \ell_2 \dots (\ell_n)^{\omega}$  with  $\forall i < n, \ell_i \oplus \ell_{i+1} \neq \emptyset$  (terminal stuttering words)

All other accepted words can be generated by duplicating letters in the above words.

Since we have only touched stuttering transitions, it is clear that the non-stuttering words of  $\mathscr{L}(\mathcal{T}'')$  are the non-stuttering words of  $\mathscr{L}(\mathcal{T}''')$ .

We now consider the case of a terminal stuttering word  $\sigma = \ell_0 \ell_1 \ell_2 \dots (\ell_n)^{\omega}$  with  $\forall i < n, \ell_i \oplus \ell_{i+1} \neq \emptyset$ .

 $(\mathcal{T}''' \subseteq \mathcal{T}^{\dagger})$  The path  $\pi'''$  that recognizes  $\sigma$  in  $\mathcal{T}'''$  has the form  $(q_0, \ell_0 \oplus \ell_1, F_0, q_1)$  $(q_1, \ell_1 \oplus \ell_2, F_1, q_2) \dots (q_n, \emptyset, \mathcal{F}, q_n)^{\omega}$  where all transitions are necessarily from  $\mathcal{T}^{\dagger}$  because we have only added in  $\mathcal{T}'''$  transitions of the form  $(q, \emptyset, \emptyset, q)$ .  $\pi'''$  is thus also an accepting path of  $\mathcal{T}^{\dagger}$  and  $\sigma \in \mathscr{L}(\mathcal{T}^{\dagger})$ .

 $(\mathcal{T}''' \supseteq \mathcal{T}^{\dagger})$  The path  $\pi^{\dagger}$  that recognizes  $\sigma$  in  $\mathcal{T}^{\dagger}$  does only stutter after  $\ell_n$ . Because this is an accepting path, it has a lasso-shape (i.e., a finite path starting from an initial state with a cycle at the end), where the cyclic part is only stuttering and accepting. Let us denote it  $\pi^{\dagger} = (q_0, \ell_0 \oplus \ell_1, F_0, q_1)(q_1, \ell_1 \oplus \ell_2, F_1, q_2) \dots (q_{n-1}, \ell_{n-1} \oplus \ell_n, F_{n-1}, q_n)(q_n, \emptyset, F_n, q_{n+1}) \dots [(q_m, \emptyset, F_m, q_{m+1}) \dots (q_l, \emptyset, F_l, q_m)]^{\omega}$ , with  $\forall i < n, \ell_i \oplus \ell_{i+1} \neq \emptyset$ .

Thanks to property 9.1, the accepting cycle  $[(q_m, \emptyset, F_m, q_{m+1}) \dots (q_l, \emptyset, F_l, q_m)]$  of  $\pi^{\dagger}$  can be replaced by an accepting self-loop  $(q_m, \emptyset, \mathcal{F}, q_m)$ . And thanks to property 9.2, the transitions from  $q_{n-1}$  to  $q_m$  can be replaced by a single transition  $(q_{n-1}, \ell_{n-1} \oplus \ell_n, F_{n-1}, q_m)$ . The resulting path  $\pi''' = (q_0, \ell_0 \oplus \ell_1, F_0, q_1)(q_1, \ell_1 \oplus \ell_2, F_1, q_2) \dots (q_{n-1}, \ell_{n-1} \oplus \ell_n, F_{n-1}, q_m)(q_m, \emptyset, \mathcal{F}, q_m)^{\omega}$  is an accepting path of  $\mathcal{T}'''$  that accepts  $\sigma$ , so  $\sigma \in \mathscr{L}(\mathcal{T}''')$ .

4. This is a classical optimization on Büchi automata.

#### 5.4 Explicit Model checking using TGTA

As for the other variants of  $\omega$ -automata, the automata-theoretic approach using TGTA has two important operations: the construction of a TGTA  $\mathcal{T}$  recognizing the negation of the stutter-invariant LTL property  $\varphi$  and the emptiness check of the product ( $\mathcal{K} \otimes \mathcal{T}$ ) of the Kripke structure  $\mathcal{K}$  with  $\mathcal{T}$ . Currently, the TGTA  $\mathcal{T}$  is built from a TGBA obtained from the translation of  $\varphi$ . In future work we plan to implement a direct translation from LTL\X to TGTA, but the construction presented above is enough to show the benefits of using TGTAs, and makes it easier to understand how TGTAs relates to TGBAs.

#### 5.4.1 Synchronous Product of a TGTA with a Kripke Structure

The product of a TGTA  $\mathcal{T}$  with a Kripke structure  $\mathcal{K}$  is an  $\emptyset$ -TGTA ( $\mathcal{K} \otimes \mathcal{T}$ ) whose language is the intersection of both languages, i.e,  $\mathscr{L}(\mathcal{K} \otimes \mathcal{T}) = \mathscr{L}(\mathcal{K}) \cap \mathscr{L}(\mathcal{T})$ .

Comparing this definition with the previous two products (for TGBA and TA) shows the double inheritance of TGTA. This product is similar to the product between a TA and a Kripke structure (Definition 25 page 47), except that it does not deal with livelock acceptance states and implicit stuttering. It is also similar to the product of a TGBA with a Kripke structure (Definition 22 page 31), except for the use of changesets on transitions, and the initial labels (U).

**Definition 31.** For a Kripke structure  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$  and a TGTA  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F} \rangle$ , the product  $\mathcal{K} \otimes \mathcal{T}$  is a  $\emptyset$ -TGTA  $\langle S_{\otimes}, I_{\otimes}, U_{\otimes}, \delta_{\otimes}, \mathcal{F}_{\otimes} \rangle$  where

- $S_{\otimes} = S \times Q$ ,
- $I_{\otimes} = \{(s,q) \in \mathcal{S}_0 \times I \mid l(s) \in U(q)\},\$
- $\forall (s,q) \in I_{\otimes}, U_{\otimes}((s,q)) = \{l(s)\},\$

96

- $\delta_{\otimes} = \{((s,q),k,F,(s',q')) \mid (s,s') \in \mathcal{R}, (q,k,F,q') \in \delta, k = (l(s) \oplus l(s'))\},\$
- $\mathcal{F}_{\otimes} = \mathcal{F}$ .

**Property 10.** We have  $\mathscr{L}(\mathfrak{K} \otimes \mathcal{T}) = \mathscr{L}(\mathfrak{K}) \cap \mathscr{L}(\mathcal{T})$  by construction.



(c) Synchronous Product  $\mathcal{K} \otimes \mathcal{T}$ 

Figure 5.6: Example of a Synchronous Product  $\mathcal{K} \otimes \mathcal{T}$  between a Kripke structure  $\mathcal{K}$  and a TGTA  $\mathcal{T}$  recognizing the LTL formula FG *p*, with acceptance conditions indicated by the black dot  $\bullet$ .

Figure 5.6 shows an example of a Synchronous Product between a Kripke structure  $\mathcal{K}$  and a TGTA  $\mathcal{T}$  recognizing the LTL formula FG p. Each state of  $\mathcal{K}$  is numbered and labeled with the set of atomic propositions (of  $\Sigma = \{p\}$ ) that hold in this state. In the TGTA representing the product  $\mathcal{K} \otimes \mathcal{T}$ , the states are labeled with pairs of the form (s,q) where  $s \in \mathcal{K}$  and  $q \in \mathcal{T}$ .

In this example, we can notice that this product using TGTA is smaller than the product using Büchi automata presented in Figure 2.12 (page 32), i.e., the product between the same Kripke structure  $\mathcal{K}$  and the TGBA  $\mathcal{A}$  recognizing the LTL formula FG p. Indeed, the synchronization of the TGTA  $\mathcal{T}$  with the stuttering parts of  $\mathcal{K}$  produces a smaller product than in the case of TGBA  $\mathcal{A}$  (compare the (sub-)product of  $\mathcal{T}$  and  $\mathcal{A}$  with the stuttering cycle of  $\mathcal{K}$ , i.e, the cycle between the states 2 and 3).

#### **5.4.2** Emptiness check (the same as TGBA)

Since a product of a TGTA with a Kripke structure is an  $\emptyset$ -TGTA, we only need an emptiness check algorithm for an  $\emptyset$ -TGTA automaton. An  $\emptyset$ -TGTA can be seen as a TGBA whose transitions are labeled by changesets instead of valuations of atomic propositions. When checking a TGBA for

#### Chapter 5. Transition-based Generalized Testing Automata (TGTA): A Single-pass and 98 Generalized New Automata

emptiness, we are looking for an accepting cycle that is reachable from an initial state. When checking an 0-TGTA for emptiness, we are looking exactly for the same thing. Therefore, because emptiness check algorithms do not look at transitions labels, the same emptiness check algorithm used for the product using TGBA (Algorithm 1) can also be used for the product using TGTA.

This is a nice feature of TGTA, not only because it gives us a one-pass emptiness check, but also because it eases the implementation of the TGTA approach in our tool Spot or in any other TGBA-based model checker. We only need to implement the conversion of TGBA to TGTA and the product between a TGTA and a Kripke structure. We discuss our implementation in the next section.

### 5.5 Experimental evaluation of TGTA

In order to evaluate the TGTA approach against the TGBA (BA) and TA approaches, an experimentation was conducted under the same conditions as Section 3.6, i.e., within the same CheckPN tool on top of Spot and using the same benchmark Inputs (formulas and models) used in the experimental comparison of BA, TGBA, TA and STA, see Section 3.6.2 page 53.

#### 5.5.1 Implementation



Figure 5.7: The experiment's architecture in SPOT. Three command-line switches control which one of the approaches is used to verify an LTL formula on a Kripke structure. The new components required by the TGTA approach are outlined in Gray.

Figure 5.7 shows the building blocks we used to implement the three approaches. The automaton used to represent the property to check has to be synchronized with a Kripke structure representing the model. Depending on the kind of automaton (TGBA, BA, TA, TGTA), this synchronous product is implemented differently. Only the TGBA and BA approaches can share the same product implementation. Approaches TA and TGTA require a dedicated product computation.

The TGBA, BA, and TGTA approaches share the same emptiness check, while a dedicated algorithm is required by the TA approach. In Figure 5.7, no direct translation is provided from LTL to TGTA (this is also true for BA and TA). This could be investigated in future work, the need being, so far, to assess their effectiveness before optimizing the translation process.

The time spent doing the conversion from LTL to TGBA and then to TGTA (bisimulation included) is measured in the benchmark of the next section (see tables column " $T_{\phi}$ " in centiseconds). This translation process is almost instantaneous, and even if its runtime could be improved (for instance with a direct translation from LTL to TGTA) it is clearly a non significant part of the run time of the different model checking approaches, where all the time is spent performing the emptiness check of the product (built on-the-fly) between the Kripke structure and the property automaton.

#### 5.5.2 Results

Table 5.1 and Table 5.2 shows how for TGBA, TA and TGTA approaches deal with toy models and random formulas. We omit data for BA since they are always outperformed by TGBA. Table 5.5 and Table 5.6 show toy models against weak-fairness formulas.

Table 5.3 and Table 5.4 show the results of the two cases studies against random, weak-fairness, and dedicated formulas issued from the studies.

These tables separate cases where formulas are verified from cases where they are violated. In the former (Tables 5.1, 5.5 and 5.3), no counterexample are found and the full state-space had to be explored; in the latter (Tables 5.2, 5.6 and 5.4) the on-the-fly exploration of the state-space stopped as soon as the existence of a counterexample could be computed.

All values shown in all tables are averaged over 100 different formulas. Indeed, we checked each model against 100 random and 100 weak-fairness formulas that had no counterexample, and against 100 random and 100 weak-fairness formulas that had a counterexample. The average and maximum are computed separately on each model against each set of formulas.

Column-wise, these tables show the average and maximum sizes (states and transitions) of: (1) the automata  $A_{\neg \varphi_i}$  expressing the properties  $\varphi_i$ ; (2) the products  $A_{\neg \varphi_i} \otimes \mathcal{K}$  of the property with the model; and (3) the subset of this product that was actually explored by the emptiness check. The emptiness check values show a third column labeled "T": this is the time (in hundredth of seconds, a.k.a. centiseconds) spent doing that emptiness check, including the on-the-fly computation of the subset of the product that is explored. In the same way, the column "T $_{\varphi}$ " shows the time (in centiseconds) spent constructing the property automata  $A_{\neg \varphi_i}$  from the formulas (this time is negligible compared to that of the emptiness check).

Figure 5.8 compares the number of visited transitions when running the emptiness check; plotting TGTA against TA and TGBA. This gives an idea of their relative performance. Each point corresponds to one of the 5600 evaluated formulas (2800 violated with counterexample as black circles, and 2800 verified having no counterexample as green crosses). Each point below the diagonal is in favor of TGTA while others are in favor of the other approach. Axes are displayed using a logarithmic scale. No comparison is presented with BA since they are less efficient than TGBA, according to the experimental evaluation presented in section 3.6.3 page 54.

All these experiments were run on a 64bit Linux system running on an Intel(R) 64-bit Xeon(R) @2.00GHz, with 10GB of RAM.

Chapter 5. Transition-based Generalized Testing Automata (TGTA): A Single-pass and 100 Generalized New Automata

			Verified properties (no counterexample)									
			1	Automator	n	Full	product	Eı	nptiness check			
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т		
	TCBA	avg	12	460	1	2753024	14166594	2753024	14166594	18656		
<b>n5</b>	TODA	max	54	2816	6	11010896	108 553 439	11010896	108 553 439	80455		
[OS]	ТЛ	avg	83	4324	12	2600539	10293714	4498653	17832896	30334		
ete	IA	max	509	37096	240	9384663	42614845	18769326	85229690	122867		
P	тста	avg	67	4584	14	2474307	9810328	2474307	9810328	16630		
	IUIA	max	349	49259	290	9596713	43944762	9596713	43944762	62911		
	TODA	avg	12	411	0	1 3 9 8 0 2 1	10139043	1 398 021	10139043	2325		
1	IGBA	max	49	3788	4	4269384	50547840	4269384	50547840	10181		
1g6	<b>T</b> A	avg	69	3118	9	1097667	6606130	2155839	12981954	3744		
Riı	IA	max	509	25773	150	2851152	19625864	5702304	39251728	11149		
	тота	avg	57	3283	10	1075198	6646513	1075198	6646513	1884		
	IGIA	max	349	31401	182	2853568	20290856	2853568	20290856	5247		
	тсра	avg	9	245	0	2038054	15037218	2038054	15037218	3017		
	IUDA	max	28	2612	4	9132417	89 397 363	9132417	89397363	15594		
IS5	TA	avg	48	1643	5	1401286	11000449	1 4 9 6 8 8 2	11732337	3084		
FI	IA	max	282	11007	98	6109887	54324903	6368802	54324903	14894		
	тста	avg	38	1713	6	1338609	10551757	1338609	10551757	2621		
	IUIA	max	247	13150	125	6109887	54324903	6109887	54324903	12570		
	TCDA	avg	8	166	0	3356053	32800737	3356053	32800737	5253		
<u>5</u>	IGBA	max	60	1994	4	20253072	258315134	20253072	258315134	44 28 1		
ban	<b>T</b> A	avg	44	1659	3	2433853	23 389 805	2433853	23 389 805	4871		
an	IA	max	277	21251	80	15272712	161 364 553	15272712	161 364 553	33129		
¥	тста	avg	36	1779	4	2285091	22038994	2285091	22038994	4385		
	IUIA	max	210	21896	89	14059472	149569931	14059472	149569931	34534		
	тсра	avg	14	636	1	4668178	27 564 474	4668178	27 564 474	9352		
0	TODA	max	81	5397	6	17947837	119545256	17947837	119545256	40273		
lol	ТЛ	avg	71	4092	20	2334154	19893200	2334154	19893200	7742		
Phi	IA	max	412	55321	235	8378151	74240975	8378151	74240975	27233		
	тста	avg	55	4221	23	2047915	17440829	2047915	17440829	6650		
	IUIA	max	269	49089	256	7381980	61 901 565	7 381 980	61901565	24212		
	TGPA	avg	15	664	0	2682881	30981263	2682881	30981263	9192		
N.	TODA	max	82	5792	4	13129644	225905117	13129644	225905117	62328		
in1	ТΛ	avg	95	6048	15	2027251	17168855	3928111	33 295 175	12207		
Sob	IA	max	481	43547	149	8169472	68645888	16338944	137 291 776	48584		
	тста	avg	85	6725	17	2042586	17645326	2042586	17645326	6 3 7 0		
		max	414	46843	137	8097792	70340609	8097792	70340609	25213		

Table 5.1: Comparison of the three approaches on toy examples with **random formulæ**, when counterexamples do not exist.
					1	/iolated prope	erties (a counter	example exis	ts)	
				Automator	n	Full	product	E	mptiness check	
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т
	тсра	avg	14	646	0	8118953	42865329	600 828	2325266	3892
15	TODA	max	47	3 3 9 2	5	50297469	364 396 821	3632979	22095045	24665
losi	ТЛ	avg	90	5683	15	8 592 222	33972863	546562	1886029	3683
ete	IA	max	656	62947	134	56166663	224108070	3178506	11478151	22221
	тста	avg	68	5389	16	7915691	31411013	535126	1850064	3540
	IUIA	max	321	40978	142	55 395 527	221 326 872	3 3 8 5 8 5 9	12455412	22971
	TGRA	avg	17	713	1	3 1 6 3 7 6 3	26260802	854229	5805080	1 3 2 9
	TODA	max	51	3788	4	9855589	108310156	5914415	65020741	13999
ng6	ТА	avg	125	7462	15	2846468	18131253	646 844	$\mathbf{3858721}$	1 1 98
Ri	17	max	472	41068	162	8960960	62840996	4923301	36659892	10659
	тста	avg	99	7247	15	2605476	17026945	640074	3911536	1085
	1017	max	383	39636	156	8712424	65764753	5438836	40403912	10615
	тсра	avg	14	496	1	10174866	95061682	1602719	10636471	2148
	IUDA	max	96	4200	3	40672843	404 072 123	11075871	110774947	18644
1S5	ТА	avg	90	4112	9	8418689	70784299	1 0 4 4 8 9 6	7693244	1968
E	IA	max	390	22702	78	31864749	267656316	8518894	63178516	16437
	тста	avg	72	4005	10	7629575	64829082	934129	6888487	1709
	IUIA	max	325	23548	82	29733270	247 258 461	7 193 693	60083469	14726
	тера	avg	13	355	0	10952627	119272618	1153850	8728297	1473
2	IUDA	max	87	2184	3	43038083	587 543 197	13023468	185410581	26750
bar	ТЛ	avg	75	2870	6	8429948	83120910	670172	5871206	1260
(an	IA	max	473	30375	94	31602068	331 833 407	7673335	77363942	15562
1	TGTA	avg	63	2986	7	7919885	79262186	631528	5523476	1113
	1017	max	321	25589	104	31771853	346223237	7673335	77364390	15411
	TGRA	avg	13	590	0	19118012	195258228	1518296	7522623	2651
	TODA	max	74	5928	3	84068722	1377479362	9032250	53881112	15880
101	ТА	avg	92	5690	14	16012398	148019289	900960	6499587	2558
Phi	IA	max	356	46498	164	66 53 5 3 2 2	696331784	7342016	64470840	23 509
.	тста	avg	76	5797	16	14738127	138827585	804739	5945101	2273
	1017	max	352	37 193	147	66964226	712347867	6062800	52054854	18939
	TGPA	avg	21	965	1	5514573	67871153	1676314	18343280	5475
N	TODA	max	74	5928	7	15127431	281 880 971	13627374	191356248	52962
lii	ТА	avg	164	11366	21	4869334	41 343 229	1248873	10418682	4110
Rob	IA	max	854	62419	115	11791872	104488704	7767216	74241202	27244
	TGTA	avg	133	11564	22	4637571	40182518	1283417	10882928	3999
	IUIA	max	485	77546	140	11726848	97390592	8539312	82390194	27829

Table 5.2: Comparison of the three approaches on toy examples with **random formulæ**, when **counterexamples exist.** 

Chapter 5.	<b>Transition-based Generalized</b>	Testing Automata (TGTA): A Single-pass an	d
102		Generalized New Automat	a

							Verified pro	perties (no cou	nterexample)		
			ĺ	1	Automator	1	Full	product	Er	nptiness check	
				st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т
		тсра	avg	16	715	1	109158	328976	109158	328976	601
		TODA	max	136	5792	4	373 039	2138028	373 039	2138028	2207
	R	ТΛ	avg	94	5768	17	94283	205464	158401	346203	867
	R	IA	max	481	43547	183	261318	604061	522636	1 208 122	2753
		тста	avg	78	5890	19	97432	$\boldsymbol{213975}$	97 432	213975	<b>530</b>
3/2		IUIA	max	414	46843	241	296747	677 842	296747	677 842	1587
33		TGRA	avg	5	72	0	96435	277426	96435	277426	524
<b>R</b>		TODA	max	18	337	2	368 898	1450005	368 898	1450005	2049
lyC	Fair	ТΔ	avg	79	1106	1	120479	262456	121096	263776	652
$\mathbf{P}_{0}$	X	IA	max	351	5920	8	633 208	1408943	633 208	1 408 943	3 3 5 9
		TGTA	avg	34	556	1	88 010	191692	88010	191692	475
		101/1	max	120	2591	4	345486	798084	345 486	798084	1828
		TGBA	-	7	576	1	345 241	760491	345241	760491	1688
	$\Phi_{1}$	TA	-	80	14590	8	342613	742815	685226	1485630	3374
		TGTA	-	79	17153	12	345 277	753798	345 277	753798	1697
		тера	avg	10	333	1	2808973	25212509	2808973	25212509	5 3 2 5
		IUDA	max	64	3813	5	13636352	147 555 158	13636352	147 555 158	28979
	Ð	T۸	avg	50	2239	9	1724291	20203618	1724291	20203618	4852
	R	IA	max	336	22614	149	8469258	108847708	8469258	108847708	25029
		тста	avg	40	2364	10	1574850	18534045	1574850	18534045	4289
		IUIA	max	279	25085	180	8469258	108847708	8469258	108847708	27737
K 8		TGRA	avg	5	39	0	3898645	34822961	3898645	34822961	7 3 4 4
AP		TODA	max	21	198	2	14452198	162156912	14452198	162156912	31117
Σ	Fair	Т۸	avg	44	407	1	2112810	24492299	2112810	24492299	6126
	X	IA	max	171	1920	3	6110748	75624744	6110748	75624744	20653
		тста	avg	18	184	0	1871267	21752673	1871267	21752673	5258
		IUIA	max	68	915	2	6110643	75624228	6110643	75624228	19865
		TGBA	-	6	165	1	46494	302350	46494	302350	50
	$ \Phi $	TA	-	9	293	2	33376	289235	33 376	289235	50
		TGTA	_	8	452	1	33376	289235	33 376	289235	49

Table 5.3: Comparison of the three approaches for the **case studies when counterexamples do not exist.** 

			[			1	Violated prope	erties (a counter	example exist	ts)	
				1	Automator	1	Full	product	Er	nptiness check	
				st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т
		тсра	avg	14	760	1	132232	380305	53901	128960	291
		IUDA	max	65	7004	9	1 295 661	5045814	192692	514905	1049
	Ð	Т۸	avg	81	6223	29	135237	296081	56518	126300	311
	$\mathbb{Z}$	IA	max	540	59971	574	1 377 784	2976709	243950	547459	1284
		тста	avg	68	6815	33	131544	$\mathbf{289575}$	52 399	116656	289
32		IUIA	max	368	67215	821	1106728	2391169	250279	561 571	1351
33		TGRA	avg	5	66	0	96849	258704	69053	177 391	372
R		TODA	max	22	516	2	432240	1584372	193264	1 0 3 1 8 9 8	1110
NO	Fair	Т۸	avg	82	1157	1	201 879	434747	113705	250573	615
Po	$\mathbb{R}$	IA	max	262	4636	5	560685	1 194 198	254328	594751	1358
		тста	avg	32	520	1	91 642	197811	66987	147418	361
		IUIA	max	142	3266	5	393 006	836499	153963	372478	840
		TGRA	avg	13	513	1	27266223	405 363 157	1850695	12575125	2702
		TODA	max	59	4298	2	91059214	1963331216	13284206	111916150	21771
	Ð	Т۸	avg	88	4818	13	22458466	298950743	965 567	9245311	2392
	R	IA	max	334	35401	165	90049281	1 300 904 178	5241327	61655512	14227
		тста	avg	71	4822	14	19806355	264723502	954 582	9085179	2243
		IUIA	max	283	38342	196	52079530	722041705	5265831	61724760	13674
<b>8</b>		TGRA	avg	6	81	0	20904101	300416331	1962765	13896236	3129
AP		TODA	max	23	414	2	101813441	1681905871	9463562	72266873	16928
Z	Fair	Т۸	avg	88	1212	1	27 53 5941	358622897	1112324	11462361	3002
	$\mathbb{R}$	IA	max	245	3749	5	93441682	1258076980	5203746	59642346	16992
		тст	avg	38	598	1	16051206	211401085	964009	9865703	2462
		IUIA	max	106	2437	3	91711581	1 302 884 523	4279478	53059472	11687

Table 5.4: Comparison of the three approaches for the **case studies when counterexamples exist.** 

Chapter 5. Transition-based Generalized Testing Automata (TGTA): A Single-pass and 104 Generalized New Automata

						Verified pr	operties (no co	unterexample)	1	
			A	Automaton		Full	product	Eı	nptiness check	
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т
	TCDA	avg	4	52	0	3071854	15745483	3071854	15745483	21163
5		max	17	357	2	9551018	87453872	9551018	87453872	70107
LSOI	T.4	avg	50	564	1	3018533	11771410	3436223	13401691	23 597
etei		max	184	3 3 6 2	4	10750205	42 228 050	11867218	47116180	78154
đ	тста	avg	25	368	1	2576578	10081563	2576578	10081563	17572
	IUIA	max	127	2998	3	6567277	26641077	6567277	26641077	43486
	TODA	avg	5	64	0	1121590	7920270	1 1 2 1 5 9 0	7920270	1903
	TGBA	max	25	357	2	3873576	38047692	3873576	38047692	7910
1g6	<b>T</b> 4	avg	58	682	1	1118511	6494556	1774991	10328239	3046
Rir	IA	max	183	3 3 5 3	3	3383372	22481284	6766744	44962568	11037
	TOTA	avg	31	476	1	878 913	5232410	878 913	5232410	1567
	IGIA	max	127	2998	3	2700824	18878992	2700824	18878992	4862
	TCDA	avg	4	41	0	3030750	23 589 355	3030750	23 589 355	4674
		max	12	212	2	11882973	114787553	11882973	114787553	23008
IS5	T.4	avg	48	515	1	2259230	17600147	2259230	17600147	4589
FI		max	171	4256	3	9764223	79707273	9764223	79707273	18521
	тста	avg	20	258	1	1908957	15050447	1908957	15050447	3711
	IUIA	max	71	1 2 9 5	3	7100604	58423901	7 100 604	58423901	13774
	TODA	avg	4	33	1	3098615	29435137	3098615	29435137	4712
Ŋ	IGBA	max	21	168	2	13014212	187004177	13014212	187004177	26143
ban	<b>T</b> A	avg	36	306	1	2083777	19156589	2083777	19156589	4112
an		max	140	2049	3	11039112	106534190	11039112	106534190	23453
X	тста	avg	16	182	1	1876505	17440799	1876505	17440799	3512
	IUIA	max	60	1185	2	10017560	98020083	10017560	98020083	20406
	тера	avg	4	35	0	4255979	26793678	4255979	26793678	9064
0		max	21	216	2	12974557	112561242	12974557	112 561 242	39806
<b>lo1</b>	TA	avg	47	501	1	3 2 2 9 6 9 0	27631036	3 2 2 9 6 9 0	27631036	10439
Phi	IA	max	289	8199	6	10987384	94141317	10987384	94141317	38378
	тста	avg	18	216	0	2278597	19552345	2278597	19552345	7184
	IUIA	max	85	2393	3	9026746	81612913	9026746	81612913	28914
	тсра	avg	5	59	0	1782133	17996764	1782133	17996764	5749
S	TODA	max	16	337	1	5029881	70344648	5029881	70344648	19263
in1	ТА	avg	54	618	1	1508082	12171559	1802576	14 591 484	5731
Sob		max	239	3706	4	3944448	32391168	7 098 368	61216768	22074
	тст	avg	28	420	1	1363362	11099030	1363362	11099030	4345
		max	120	2591	4	3649536	30605312	3649536	30605312	11669

Table 5.5: Comparison of the three approaches on toy examples with **weak-fairness formulæ**, **when counterexamples do not exist.** 

					/	/iolated prope	rties (a counter	example exist	s)	
			1	Automator	ı	Full p	product	En	nptiness check	
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т
	тсра	avg	5	78	0	3731947	17751951	522414	1759041	3438
15	TODA	max	18	351	2	13 508 730	79589415	3696481	16733420	24769
losi	ТЛ	avg	95	1400	2	10950499	42790086	641997	2123401	4370
ete	IA	max	350	5968	13	30354888	120985972	7637347	29621420	51473
	тота	avg	36	598	1	3563419	14000819	506574	1641570	3400
	IUIA	max	139	2666	4	11920402	47 292 924	3611716	14056752	24606
	TGBA	avg	7	123	1	2000787	15868857	890567	5736568	1410
	TODA	max	38	1168	3	4670124	53246768	2766283	26917302	6101
ng6	ТА	avg	128	2009	2	3456785	21069783	905265	5156327	1621
Ri	IA	max	388	12783	14	6152232	43283696	2747578	17464827	4793
	тста	avg	55	997	1	1640027	10134116	677559	3974017	1183
	IUIA	max	220	7174	8	4234800	28442432	2317315	14895284	4336
	тсра	avg	5	68	0	7978129	75701408	1372877	7544397	1607
	IUDA	max	23	370	2	27 570 312	293219919	6676584	53451994	9961
1S5	TA	avg	71	871	1	10007988	81736556	848557	5707348	1 5 2 2
E	IA	max	245	3766	5	27733464	249573778	6425987	51962113	12307
	тста	avg	33	474	1	5682339	46813094	709810	4738120	1223
	IUIA	max	131	2214	3	19497180	165240528	5062934	41766785	10619
	тера	avg	5	66	0	7040236	73867374	1344870	9279129	1562
2	IUDA	max	17	305	2	25423161	388961954	7950633	80041971	11336
bai	ТА	avg	69	849	1	8 102 548	77515606	809467	6987095	1474
an	IA	max	237	3655	3	32732392	329125090	5 501 242	53375472	10218
<b></b>	тста	avg	31	459	1	4837392	46966950	718473	6247954	1262
	IUIA	max	128	2446	3	21 299 320	218268197	4580718	46290149	8953
	TGRA	avg	5	69	0	10515366	102622810	833663	3762217	1357
	TODA	max	25	392	3	24496745	402197001	5017174	32273634	9802
<b>I</b> 01	ТА	avg	81	1027	1	16779360	150030607	742819	5019286	1943
Phi	IA	max	267	4339	5	43186049	394172486	6703064	55546036	20284
.	тста	avg	35	525	1	7800504	70806206	442830	2727715	1103
	IUIA	max	116	3004	4	20986825	203614146	3405659	28381991	9712
	TGRA	avg	9	157	1	4006015	46910526	1486839	14360747	4691
N	TODA	max	44	1168	2	14802883	195152925	5666009	78112291	21619
lii	ТА	avg	143	2467	2	5832674	48 201 707	1 348 406	10788993	4327
Rob	IA	max	466	12783	14	18204684	151155748	6486142	53115008	18844
	TGTA	avg	65	1272	1	3428985	28693985	1188348	9648116	3825
	IUIA	max	228	7174	7	12249088	102016000	4618329	38837339	13884

Table 5.6: Comparison of the three approaches on toy examples with **weak-fairness formulæ**, **when counterexamples exist.** 

Chapter 5. Transition-based Generalized Testing Automata (TGTA): A Single-pass and 106 Generalized New Automata



(b) TGTA against TGBA approaches

Figure 5.8: Performance (number of transitions explored by the emptiness check) of TGTA against TA and TGBA.

#### 5.5.3 Discussion

**On verified properties** the results are very straightforward to interpret when looking at the number of transitions explored by the emptiness check. TA outperform TGBA except for both Random and weak-fairness properties against Peterson, Ring, Robin and PolyORB. These are typical cases where the TA emptiness check has to perform two passes: this can be observed in the tables 5.1, 5.3 and 5.5 when the number of transitions visited by the emptiness check is on the average twice the number of transitions of the product.

In these three cases, the TGTA approach, with its single-pass emptiness check, is a clear improvement over TA. On the scatter plots of Figure 5.8a, these cases where the TGTA approach is twice faster than TA's, appear as a linear cloud of green crosses below the diagonal (we recall that the axes are displayed using a logarithmic scale).

In the other cases where TA need only one pass on the average (e.g. Kanban, MAPK, Philo), TGTA and TA have similar performance, with a slight advantage for TGTA because the products are smaller, especially for weak-fairness formulas because TGTA represent more concisely this kind of formulas using a large number of generalized acceptance conditions (similar to TGBA).

To summarize, the TGTA approach outperforms TGBA and TA approaches in all cases on verified properties.

**On violated properties**, we recall that it is difficult to interpret the scatter plots of results because the emptiness check is an on-the-fly algorithm. It stops as soon as it finds a counterexample. Thus, the exploration order of non-deterministic transitions of TGBA, TA and TGTA changes the number of states and transitions to be explored in the product before a counterexample is found.

However, if we analyze more precisely tables 5.2, 5.4 and 5.6, we observe that the TGTA approach produces the smallest products on the average. This allows the TGTA approach to seek a counterexample in a smaller product and therefore have a better chance to find it faster. Thus, we observe that in the majority of cases the emptiness check of TGTA approach explores less states and transitions on the average than TGBA or TA.

In the tables of results, we generally observe for both verified and violated properties that:

- Although the TGTA constructed from properties are usually larger than TGBA (and even larger than BA), the average sizes of the products in TGTA approach are smaller than the average sizes of the products in TGBA approach (and BA approach, see section 3.6.3 page 54). We believe this is due to the elimination of useless stuttering-transitions in TGTA (see Section 5.3.2). specificity.
- In addition, if we compare the automata sizes for TA versus TGTA, we observe that TGTA are smaller than TA in all tables in terms of average numbers of states, and in terms of average numbers of states and transitions for weak-fairness formulas (in Tables 5.2, 5.6, 5.3 and 5.4). We believe this is due to the fact that TGTA represent more concisely the LTL formulas using (multiple) generalized acceptance conditions, especially for weak-fairness formulas (for which the number of acceptance conditions is greater, in our experiments we have |F| = 3.51 on the average for weak-fairness formulas, while |F| = 1.32 on the average for random formulas). This is the consequence of the fact that a TA is built from a BA while a TGTA is built from a TGBA. Thus, TGTA can take advantage from the fact that

### Chapter 5. Transition-based Generalized Testing Automata (TGTA): A Single-pass and 108 Generalized New Automata

TGBA are more concise than BA. TGBA are smaller [49, 36] because they use generalized acceptance conditions. For example, the BA of  $\varphi = GFa \wedge GFb$  (Figure 2.4a) contains 3 states and it is transformed <sup>1</sup> into a TA that contains 10 states (30 transitions); the TGBA of  $\varphi = GFa \wedge GFb$  (Figure 2.7a) is only composed by one state and it is therefore transformed <sup>1</sup> into a TGTA that only contains 4 states (16 transitions).

#### 5.5.4 Experimental Results once the TGBA is improved by simulation-reduction

Recently, a new TGBA optimization was added to SPOT including the simulation-reduction [3] of TGBA. Unfortunately, we did not have the time to implement this simulation-reduction for TA and TGTA. However, since the construction of TA and TGTA depend on TGBA (we recall that TA is constructed from BA, which is obtained from TGBA by degeneralization), then the reduction of TGBA can also reduce both TA and TGTA sizes.

The tables and scatter plots presented in the Appendix A show the impact of this optimization on the experimental results presented in the previous Section 5.5.2. This impact is positive for the three approaches. Indeed, if we compare the results of the previous Section 5.5.2 against the tables of Appendix A, we observe that the simulation-reduction of TGBA also reduces the TA and TGTA sizes.

These reductions produce smaller products on average and thus improve the performance of the three approaches. However, this does not change the result of the comparison of the three approaches: for verified formulas, TGTA remains more efficient than the other approaches; for violated formulas, the results are still difficult to interpret.

## 5.6 Conclusion

In the previous chapters, we have shown that TA outperformed BA and sometimes TGBA for unverified properties (i.e., when a counterexample was found). However, this was not the case when no counterexample was computed since the entire product had to be visited twice to check for each acceptance mode of a TA (Büchi acceptance or livelock-acceptance).

In this chapter, we propose a new type of  $\omega$ -automaton for stutter-invariant properties, called Transition-based Generalized Testing Automata (TGTA).

TGTA combines advantages observed on both TA and TGBA:

- From TA, it reuses the labeling of transitions with changesets, and the elimination of the useless stuttering-transitions, but without requiring a second pass in the emptiness check of the product.
- From TGBA, it inherits the use of generalized acceptance conditions on transitions.

TGTA have been implemented in Spot easily, because only two new algorithms are required: the conversion of a TGBA into a TGTA, and a new definition of a product between a TGTA and a Kripke structure.

We have run benchmarks to compare TGTA against TA and TGBA (BA). Experiments reported that, in most cases, TGTA produce the smallest products on the average and they outper-

<sup>1</sup>http://spot.lip6.fr/ltl2tgba.html

form TA and TGBA when no counterexample is found in the system, but they are comparable when the property is violated, because in this case the on-the-fly algorithm stops as soon as it finds a counterexample without exploring the entire product.

We conclude that there is nothing to lose by using TGTA to verify stuttering-insensitive properties, since they are always at least as good as TA and TGBA.

We believe that TGTA is better than TA because TGTA does not require a second pass during the emptiness check and because TGTA represent more concisely the LTL formulas using (multiple) generalized acceptance conditions, as observed in our experiments for weak-fairness formulas (for which the number of acceptance conditions is greater than random formulas).

Compared to TGBA, we believe that TGTA is better thanks to the elimination of the useless stuttering-transitions during the TGTA construction (Section 5.3.2). This elimination exploits the fact that the TGTA are specific to stutter-invariant formulas, while the TGBA does not exploit at all this specificity.

After this elimination of useless stuttering-transitions, the obtained TGTA represents all the stuttering-transitions with only self-loops on all states (see the "stuttering-normalization constraint" of the TGTA Definition 30). This advantage of TGTA will be more clearly exploited in the symbolic approach presented in the next chapter. This symbolic model checking approach using TGTA allows us to tackle much larger state-spaces than in explicit model checking.

## Part III

# Using TGTA to improve Symbolic/Hybrid Model Checking

## CHAPTER 6

# Symbolic LTL Model Checking using TGTA

#### Contents

6.1	Introd	luction
6.2	Symbo	olic LTL Model Checking
	6.2.1	Symbolic Kripke Structure
	6.2.2	Symbolic Büchi Automata TGBA 115
	6.2.3	Symbolic Product of a TGBA with a Kripke structure
	6.2.4	Symbolic Emptiness Check algorithm
6.3	TGTA	-based Symbolic LTL Model Checking 117
	6.3.1	Symbolic TGTA 117
	6.3.2	Naive Symbolic Product of TGTA with a Kripke structure 117
	6.3.3	Adjusting the Symbolic Transition Relation of the Kripke Structure to TGTA 118
	6.3.4	Exploiting stuttering transitions to Improve Saturation in the TGTA Approach119
6.4	Exper	imental evaluation
	6.4.1	Implementation
	6.4.2	Using ETF to build the transition relation of a changeset-based symbolic
		Kripke structure         122
	6.4.3	Benchmark
	6.4.4	Results
6.5	Concl	usion

## 6.1 Introduction

In the previous chapter, we showed how to generalize the Testing Automata (TA) using several acceptance sets, and allowing a single-pass efficient emptiness check. Our experimental comparison showed these Transition-based Generalized Testing Automata (TGTA) to be superior to Büchi Automata in the *explicit approach* for model-checking of stutter-invariant properties. In this explicit approach, the automata and their products were represented as explicit graphs.

Another implementation of this procedure is the *symbolic approach* where the automata and their products are represented by means of decision diagrams (a concise way to represent large

sets or relations) [17]. Encoding generalized Büchi automata is pretty common [75]. With such encoding, we can compute, in one step, the sets of all direct successors (*PostImage*) or predecessors (*PreImage*) of any set of states. Using this technique, there have been a lot of propositions for symbolic emptiness-check algorithms [43, 80, 57]. These symbolic algorithms manipulate standard BFS-based fixed-points on the transition relation of the product which can be optimized using saturation techniques [20, 85].

To the best of our knowledge, these algorithms do not provide efficient optimizations specific to stutter-invariant properties, and Testing Automata have never been used in symbolic model checking. In this chapter, we propose and evaluate a symbolic approach for model checking using TGTA [7], and compare it to the symbolic approach using TGBA. In particular, we show that the computation of fixpoints on the transition relation of the product can be sped up with a dedicated evaluation of stuttering transitions. The implementation uses the saturation technique introduced by [20] that departs from standard BFS-based approaches for the symbolic fixed-point computation. Saturation nicely fits well with TGTA. Indeed, we exploit a separation of the transition relation into two terms, one of which greatly benefits from saturation techniques.

This chapter is organized as follows. Section 6.2 presents the symbolic model-checking approach for TGBA. For generality we define our symbolic structures using predicates over state variables in order to remain independent of the variant of Decision Diagrams used to actually implement the approach. Section 6.3 focuses on the encoding of TGTA in the same framework. We first show how a TGTA can be encoded, then we show how to improve the encoding of the Kripke structure and the product to benefit from saturation in the encoding of stuttering transitions in the TGTA. Finally, Section 6.4 compares the two approaches experimentally with an implementation that uses hierarchical Set Decision Diagrams (SDD) [85] (a particular type of Decision Diagrams on integer variables, on which we can apply user-defined operations). We concentrate on a comparison of TGBA versus TGTA and on the impact of the saturation technique. On our large, BEEM-based benchmark (presented in Section 6.4.3) and huge set of LTL formulas, our symbolic encoding of TGTA appears to be superior to TGBA.

## 6.2 Symbolic LTL Model Checking

We first present how to perform the automata-theoretic approach to LTL model checking using symbolic encodings of TGBA and Kripke structures. This setup will serve as a baseline to measure our improvements from later sections.

#### 6.2.1 Symbolic Kripke Structure

In symbolic model checking we encode such a Kripke structure with predicates that represent sets of states or transitions [79]. These predicates are then implemented using decision diagrams [17].

**Definition 32** (Symbolic Kripke Structure). A Kripke structure  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$  can be encoded by the following predicates where  $s, s' \in S$  and  $\ell \in \Sigma$ :

- $P_{S_0}(s)$  is true iff  $s \in S_0$ ,
- $P_{\mathcal{R}}(s,s')$  is true iff  $(s,s') \in \mathcal{R}$ ,

•  $P_l(s, \ell)$  is true iff  $l(s) = \ell$ .

In the sequel, we use the notations  $S_0(s)$ , R(s,s') and  $L(s,\ell)$  instead of  $P_{S_0}(s)$ ,  $P_{\mathcal{R}}(s,s')$  and  $P_l(s,\ell)$ . A Symbolic Kripke structure is therefore a triplet of predicates  $K = \langle S_0, R, L \rangle$  on state variables.

Variables *s* and *s'* used above are typically implemented using decision diagrams to represent either a state or a set of states. In a typical encoding [17], states are represented by conjunctions of Boolean variables. For instance if  $S = \{0,1\}^3$ , a state s = (1,0,1) would be encoded as  $s_1\bar{s}_2s_3$ . Similarly,  $s_1s_3$  would encode the set of states  $\{(1,0,1),(1,1,1)\}$ . With this encoding,  $S_0$ , *R* and *L* are propositional formulas which can be implemented with BDDs or other kind of decision diagrams. In our implementation, we used SDDs on integer variables [85].

#### 6.2.2 Symbolic Büchi Automata TGBA

We chose the TGBA to represent the negation of the LTL property to verify, because generalized acceptance is classically used in symbolic model checking [43] and using transition-based acceptance is not a problem [75]. People working with Generalized Büchi automata (GBA) can adjust to our definitions by "pushing" the acceptance of states to their outgoing transitions [36].

Like Kripke structures, a TGBA can be encoded by predicates [79] on state variables.

**Definition 33** (Symbolic TGBA). A TGBA  $\langle Q, I, \delta, \mathcal{F} \rangle$  is symbolically encoded by a the predicates  $\langle I, \Delta, \{\Delta_f\}_{f \in \mathcal{F}} \rangle$  where:

- I(q) is true iff  $q \in I$ ,
- $\Delta(q, \ell, q')$  is true iff  $(\exists F \in 2^{\mathcal{F}}, (q, \ell, F, q') \in \delta)$ ,
- For every  $f \in \mathcal{F}$ , the predicate  $\Delta_f$  is defined by:  $\Delta_f(q, \ell, q')$  is true iff  $(\exists (q, \ell, F, q') \in \delta, f \in F)$ .

#### 6.2.3 Symbolic Product of a TGBA with a Kripke structure

We now show how to build a synchronous product by composing the symbolic representations of a TGBA with that of a Kripke structure, inspired from Sebastian *et al.* [79].

**Definition 34** (Symbolic Product for TGBA). *Given a Symbolic Kripke structure*  $K = \langle S_0, R, L \rangle$ and a Symbolic TGBA  $A = \langle I, \Delta, \{\Delta_f\}_{f \in \mathcal{F}} \rangle$  sharing a set AP of atomic propositions, the Symbolic Product  $K \otimes A = \langle P_0, T, \{T_f\}_{f \in \mathcal{F}} \rangle$  is defined by the predicates  $P_0$ , T and  $T_f$  encoding respectively the set of initial states, the transition relation and the acceptance transitions of the product:

- (s,q) denotes the state variables of the product (s for the Kripke structure and q for TGBA),
- $P_0(s,q) = S_0(s) \wedge I(q)$ ,
- $T((s,q),(s',q')) = \exists \ell [R(s,s') \land L(s,\ell) \land \Delta(q,\ell,q')], \text{ where } (s',q') \text{ encodes the next state variables,}$
- $\forall f \in \mathcal{F}, T_f((s,q), (s',q')) = \exists \ell [R(s,s') \land L(s,\ell) \land \Delta_f(q,\ell,q')].$

The labels  $\ell$  are used to ensure that a transition  $(q, \ell, q')$  of A is synchronized with a state s of K such that  $L(s, \ell)$ . This way, we ensure that the product recognizes only the executions of K that are also recognized by A. However we do not need to remember how product transitions are labeled to check  $K \otimes A$  for emptiness. Therefore, a product can be seen as a TGBA without labels on transitions. This explains why the predicate T((s,q),(s',q')) does not take as an argument a variable ( $\ell$ ) to encode the labels of transitions (as in  $\Delta(q, \ell, q')$ ).

In symbolic model checkers, the exploration of the product is based on the following *PostImage* operation [79]. For any set of states encoded by a predicate *P*, *PostImage*(*P*)  $(s',q') = \exists (s,q) [P(s,q) \land T((s,q),(s',q'))]$  returns a predicate on state variables (s',q') encoding the set of states reachable in one step from the set of states encoded by *P*.

Because in TGBA the acceptance conditions are based on transitions, we also define PostImage(P, f) to compute the successors of P reached using the transitions labeled with the acceptance condition  $f \in \mathcal{F}$ :  $PostImage(P, f)(s', q') = \exists (s, q) [P(s, q) \land T_f((s, q), (s', q'))].$ 

These two operations are at the heart of the symbolic emptiness check presented in the next section.

#### 6.2.4 Symbolic Emptiness Check algorithm

```
1 Input: PostImage, P_0 and \mathcal{F}
2 begin
3
       P \leftarrow \text{Reach}(P_0)
                                                                           1 Reach(P)
4
        while P changes do
                                                                                   while P changes do
                                                                          2
             while P changes do
5
                                                                                        P \gets P \cup \texttt{PostImage}\left(P\right)
                                                                          3
                  P \leftarrow \texttt{PostImage}(P)
6
                                                                                   return P
                                                                           4
             for f in \mathcal{F} do
7
                  P \leftarrow \text{Reach}(\text{PostImage}(P, f))
8
        return P = \emptyset
9
```

Figure 6.1: Forward-variant of OWCTY, a symbolic emptiness check.

One way to check if a product is not empty is to find a reachable Strongly Connected Component that contains transitions from all acceptance sets (we call it an *accepting SCC*). Figure 6.1 shows such an algorithm implemented using symbolic operations. It mimics the algorithm FEASI-BLE of Kesten et al. [57] and can be seen as a forward variant of OWCTY (One Way Catch Them Young [43]) that uses *PostImage* computations instead of *PreImage*. Line 3 computes the set *P* of all reachable states of the product. The main loop on lines 4–8 refines *P* at each iteration. Lines 5–6 keep only the states of *P* that can be reached from a cycle in *P*. Lines 7–8 then remove all cycles that never visit some acceptance set  $f \in \mathcal{F}$ . Eventually the main loop will reach a fixpoint where *P* contains all states that are reachable from an accepting SCC. The product is empty *iff* that set is empty. There are many variants of such symbolic emptiness checks [43, 80, 57]. We selected this variant mainly for its simplicity, as our contributions are mostly independent of the chosen algorithm: essentially, we will improve the cost of computing Reach(P) (used lines 3 and 8).

## 6.3 TGTA-based Symbolic LTL Model Checking

In this section, we show how to encode a TGTA and a product for symbolic model checking using TGTA instead of TGBA.

#### 6.3.1 Symbolic TGTA

A TGTA can be encoded symbolically in a similar way as we encoded a TGBA.

**Definition 35** (Symbolic TGTA). A TGTA  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F} \rangle$  is symbolically encoded by the predicates  $\langle U_0, \Delta^{\oplus}, \{\Delta_f^{\oplus}\}_{f \in \mathcal{F}} \rangle$  where:

- $U_0(q,\ell)$  is true iff  $(q \in I) \land (U(q) = \ell)$
- $\Delta^{\oplus}(q,k,q')$  is true iff  $(\exists F \in 2^{\mathcal{F}}, (q,k,F,q') \in \delta)$ ,
- For every  $f \in \mathcal{F}$ , the predicate  $\Delta_f^{\oplus}$  is defined by:  $\Delta_f^{\oplus}(q,k,q')$  is true iff  $(\exists (q,k,F,q') \in \delta, f \in F).$

The predicates encoding the symbolic TGTA are the same as those encoding the symbolic TGBA, but they are based on changesets (encoded by the variable k) instead of valuations (variable  $\ell$ ).

The predicate  $U_0(q, \ell)$  encodes the set of initial states and their valuations. The predicate  $\Delta^{\oplus}(q, k, q')$  encodes the set of transitions of the TGTA, with the variable *k* encodes the changeset between *q* and *q'*. For each  $f \in \mathcal{F}$ , a predicate  $\Delta_f^{\oplus}(q, k, q')$  encodes the transitions labeled with the acceptance condition *f*.

#### 6.3.2 Naive Symbolic Product of TGTA with a Kripke structure

The product between a TGTA and a Kripke structure is similar to the TGBA case, except that we have to deal with changesets. The transitions (s, s') of a Kripke structure that must be synchronized with a transition (q, k, q') of a TGTA are all the transitions such that the label of *s* and *s'* differs by the changeset *k*: in the definition below, this is encoded by the terms  $L(s, \ell) \wedge L(s', \ell')$ , with  $(\ell \oplus \ell') = k$ .

**Definition 36** (Naive Symbolic Product for TGTA). *Given a Symbolic Kripke structure*  $K = \langle S_0, R, L \rangle$  and a Symbolic TGTA  $A^{\oplus} = \langle U_0, \Delta^{\oplus}, \{\Delta_f^{\oplus}\}_{f \in \mathcal{F}} \rangle$  sharing the same set of atomic propositions AP, the Symbolic Product  $K \otimes A^{\oplus} = \langle P_0, T, \{T_f\}_{f \in \mathcal{F}} \rangle$  is defined by the following predicates:

- The set of initial states is encoded by:  $P_0(s,q) = \exists \ell [S_0(s) \land L(s,\ell) \land U_0(q,\ell)].$
- A naive encoding of the transition relation of the product is:  $T(s,q,s',q') = \exists k \left[ R(s,s') \land \left( \exists \ell, \ell' [L(s,\ell) \land L(s',\ell') \land xor(\ell,k,\ell')] \right) \land \Delta^{\oplus}(q,k,q') \right], \text{ with the predicate } xor(\ell,k,\ell') \text{ is true iff } (\ell \oplus \ell') = k$

• The definition of  $T_f$  is similar to T by replacing  $\Delta^{\oplus}$  with  $\Delta_f^{\oplus}$ .

The formulas of PostImage(P) and PostImage(P, f) in TGTA approach are the same as in the TGBA approach, with the new expressions of T and  $T_f$  defined above for  $K \otimes A^{\oplus}$ .

This naive definition of *T* contains the term  $\exists \ell, \ell' [L(s,\ell) \land L(s',\ell') \land xor(\ell,k,\ell')]$  which requires several symbolic operations. In the next section, we show how to re-encode the transition relation of Kripke structures to remove this term.

#### 6.3.3 Adjusting the Symbolic Transition Relation of the Kripke Structure to TGTA

In order to reduce the number of symbolic operations in T and  $T_f$ , we introduce a changeset-based encoding of a Kripke structure (only the transition relation changes).

**Definition 37** (Changeset-based symbolic Kripke structure). *A Kripke structure*  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$ , *can be encoded by a changeset-based symbolic Kripke structure*  $K^{\oplus} = \langle S_0, R^{\oplus}, L \rangle$ , *where:* 

- the predicate  $\mathbb{R}^{\oplus}(s,k,s')$  is true iff  $((s,s') \in \mathcal{R} \land (l(s) \oplus l(s')) = k)$ ,
- the predicates S<sub>0</sub> and L have the same definition as for a symbolic Kripke structure K of Definition 32.

The changeset-based symbolic transition relation  $R^{\oplus}(s,k,s')$  of a Kripke structure is similar to the symbolic transition relation  $\Delta^{\oplus}(q,k,q')$  of a TGTA. It encodes the transitions  $((s,s') \in \mathcal{R}$  with the variable k encodes the changeset between the two valuations l(s) and l(s').

In practice, the (changeset-based or not) symbolic transition relation of a Kripke structure should be constructed directly from the model and atomic propositions of the formula to check. In Section 6.4.2, we discuss how we build such changeset-based Kripke structures in our setup.

Adjusting the symbolic encoding of the Kripke structure to TGTA, allows us to obtain the following natural definition of the symbolic product using TGTA:

**Definition 38** (Symbolic Product for TGTA). *Given a changeset-based Symbolic Kripke structure*  $K^{\oplus} = \langle S_0, R^{\oplus}, L \rangle$  and a Symbolic TGTA  $A^{\oplus} = \langle U_0, \Delta^{\oplus}, \{\Delta_f^{\oplus}\}_{f \in \mathcal{F}} \rangle$  sharing the same set of atomic propositions AP, the Symbolic Product  $K^{\oplus} \otimes A^{\oplus} = \langle P_0, T, \{T_f\}_{f \in \mathcal{F}} \rangle$  is defined by the following predicates:

- The set of initial states is encoded by:  $P_0(s,q) = \exists \ell \left[ S_0(s) \land L(s,\ell) \land U_0(q,\ell) \right]$
- The transition relation of the product is:  $T((s,q),(s',q')) = \exists k \left[ R^{\oplus}(s,k,s') \land \Delta^{\oplus}(q,k,q') \right]$
- The definition of  $T_f$  is similar to T by replacing  $\Delta^{\oplus}$  with  $\Delta_f^{\oplus}$ .

The definitions of PostImage(P) and PostImage(P, f) are the same as in the TGBA approach, with the new expressions of T and  $T_f$  above.

As for the product in TGBA approach, the product in TGTA approach is a TGBA (or a TGTA) without labels on transitions, and the same emptiness check algorithm (Figure 6.1) can be used for the two products.

## 6.3.4 Exploiting stuttering transitions to Improve Saturation in the TGTA Approach

Among symbolic approaches for evaluating a fixpoint on a transition relation, the *saturation* algorithm offers gains of one to three orders of magnitude [20] in both time and memory, especially when applied to asynchronous systems [18].

The saturation algorithm does not use a breadth-first exploration of the product, i.e., each iteration in the function Reach (Figure 6.1) is not a "global" *PostImage()* computation. Saturation instead recursively repeats "local" fixed-points by recognizing and exploiting transitions locality and identity transformations on state variables [18]. This algorithm considers that the system state consists of *n* discrete variables encoded by a Decision Diagram, and that the transition relation is expressed as a disjunction of terms called transition clusters. Each cluster typically only reads or writes a limited subset consisting of  $n' \leq n$  variables, called the *support* of the cluster. During the least fixpoint computing the reachable states, the saturation technique consists in reordering [51] the evaluation of ("local" fixed-points on) clusters in order to avoid the construction of (useless) intermediate decision diagram nodes.

The algorithm to determine an ordering for saturation is based on the support of each cluster.

We now show how to decompose the transition relation of the product  $(K^{\oplus} \otimes A^{\oplus})$  to exhibit clusters having a smaller support, favoring the saturation technique.

We base our decomposition on the fact that in a TGTA, all stuttering transitions are self-loops and every state has a stuttering self-loop (see the stuttering-normalization constraint in Definition 30 of TGTA). Therefore, stuttering transitions in the Kripke structure can be mapped to stuttering transitions in the product regardless of the TGTA state.

Let us separate stuttering and non-stuttering transitions in the transition relation *T* of the symbolic product  $K^{\oplus} \otimes A^{\oplus}$  (Definition 38):

$$T((s,q),(s',q')) = \left( R^{\oplus}(s,\emptyset,s') \land \Delta^{\oplus}(q,\emptyset,q') \right) \lor \left( \exists k \left[ R^{\oplus}_*(s,k,s') \land \Delta^{\oplus}_*(q,k,q') \right] \right)$$

where  $R^{\oplus}_*$  and  $\Delta^{\oplus}_*$  encode respectively the non-stuttering transitions of the model and of the TGTA:

- $\Delta^{\oplus}_*(q,k,q')$  is true iff  $\Delta^{\oplus}(q,k,q') \land (k \neq \emptyset)$
- $R^{\oplus}_*(s,k,s')$  is true iff  $R^{\oplus}(s,k,s') \land (k \neq \emptyset)$

On the one hand, according to Definition 35, the predicate  $\Delta^{\oplus}(q, \emptyset, q')$  is defined by the following equivalence:

$$\forall (q,q') \in Q^2 : \left[\Delta^{\oplus}(q,\emptyset,q') \Longleftrightarrow \left(\exists F \in 2^{\mathcal{F}}, (q,\emptyset,F,q') \in \delta\right)\right]$$
(6.1)

On the other hand, according to the stuttering-normalization constraint of Definition 30 of TGTA (i.e., all stuttering transitions are self-loops and every state has a stuttering self-loop):

$$\forall (q,q') \in Q^2 : \left[ \left( \exists F \in 2^{\mathcal{F}}, (q, \emptyset, F, q') \in \delta \right) \Longleftrightarrow (q = q') \right] \tag{6.2}$$

Combining the two equations (6.1) and (6.2), we obtained the following third equivalence:

$$\forall (q,q') \in Q^2 : \left[\Delta^{\oplus}(q, \emptyset, q') \Longleftrightarrow (q = q')\right]$$
(6.3)

In the other words, the predicate  $\Delta^{\oplus}(q, \emptyset, q')$  encodes the set of all self-loops of the TGTA and can be replaced by the identity predicate: equal(q, q'), simplifying *T* as:

$$T((s,q),(s',q')) = \left(\underbrace{R^{\oplus}(s,\emptyset,s') \wedge equal(q,q')}_{T_{\emptyset}((s,q),(s',q'))}\right) \vee \left(\underbrace{\exists k \left[R^{\oplus}_{*}(s,k,s') \wedge \Delta^{\oplus}_{*}(q,k,q')\right]}_{T_{*}((s,q),(s',q'))}\right) \tag{6.4}$$

The transition relation of equation (6.4) is a disjunction of  $T_*$ , synchronizing updates of both TGTA and Kripke structure, and  $T_0$ , corresponding to the stuttering transitions of the Kripke structure. Since all states in the TGTA have a stuttering self-loop,  $T_0$  does not depend on the TGTA state. In practice, the predicate equal(q,q') is an identity relation for variable q [18] and is simplified away (i.e., the term  $T_0$  can be applied to a decision diagram without consulting or updating the variable q [51]). Hence q is not part of the clusters supports in  $T_0$  (while q is part of the clusters supports in  $T_*$ ). This gives more freedom to the saturation technique for reordering the application of clusters in  $T_0$ .

Note that in the product of TGBA with Kripke structure (Definition 34) there is no  $T_0$  that could be extracted since there is no stuttering hypothesis in general. This severely limits the possibilities of the saturation algorithm in the TGBA approach.

In the symbolic emptiness check presented in Figure 6.1, the function Reach corresponds to a least fixpoint performed using saturation. As we shall see experimentally in the next section, the better encoding of  $T_{\emptyset}$  (without q in its support) in the product of TGTA with Kripke structure, greatly favors the saturation technique, leading to gains of roughly one order of magnitude.

The improvement of T proposed in this section is not applicable to  $T_f$  because the term  $\Delta_f^{\oplus}(q, \emptyset, q')$  encodes only a subset of self-loops of the TGTA (not the all self-loops as  $\Delta^{\oplus}(q, \emptyset, q')$ ), hence the expressions of  $T_f$  and *Postimage*(P, f) are the same as the previous section.

Table 6.1 summarizes all our definitions for the predicates encoding the automaton, the Kripke structure, and their product, in the four approaches presented in Sections 6.2 and 6.3.

The results of the TGTA based symbolic approach will be showed in the next section presenting our experimentation.

### 6.4 Experimental evaluation

We now compare the symbolic approaches presented in this chapter. The symbolic modelchecking approach using TGBA, presented in Section 6.2 serves as our baseline. We first describe our implementation and selected benchmarks, prior to discussing the results.

Initial states $Q_0(q)$ $U_0(q,t)$ Initial states $\Delta(q,t,q')$ $U_0(q,t)$ Initial states $\Delta(q,t,q')$ $\Delta^{\oplus}(q,t,q')$ Initial states $\forall f \in F, \Delta_{f}^{\oplus}(q,t,q')$ $\forall f \in F, \Delta_{f}^{\oplus}(q,t,q')$ Initial states $\forall f \in F, \Delta_{f}^{\oplus}(q,t,q')$ $\Delta^{\oplus}(q,t,q')$ Initial states $\forall f \in F, \Delta_{f}^{\oplus}(q,t,q')$ $\Delta^{\oplus}(q,t,q')$ Initial states $\Delta(q,t,q')$ $\Delta(r,t,q')$ Initial states $\Delta(q,t,q')$ $\Delta(r,t,q')$ Initial states $\Delta(s,q,s',q')$ $\Delta(r,t,q')$ Initial states $\Delta(s,q,s',q')$ $\Delta(r,t,q')$ Initial states $\Delta(s,q',d')$ $\Delta(r,t,s') \wedge \Delta_{\theta}(q,t,q')$ Initial states $\Delta(s,q,s',q')$ $\Delta(r,t,s') \wedge \Delta_{\theta}(q,t,q')$ Initial states $\Delta(r,q,s',q')$ $\Delta(r,t,s') \wedge \Delta_{\theta}(q,t,q')$ Initial states $\Delta(r,q,s',q')$ $\Delta(r,t,s') \wedge \Delta_{\theta}(q,t,q')$ Initial states $\Delta(r,q,s')$ $\Delta(r,t,s') \wedge \Delta_{\theta}(q,t,q')$ Initial states $\Delta(r,t,s') \wedge \Delta_{\theta}(q,t,q')$ Initial states $\Delta(r,t,$			TGBA approach (our baseline) Section 6.2	Naive TGTA approach Sections 6.3.1 and 6.3.2	First improvement of the TGTA approach Sections 6.3.1 and 6.3.3	Second improvement of the TGTA approach Sections 6.3.1, 6.3.3 and 6.3.4
Anotomat Intensitions $\Delta(q,\ell,q')$ $\Delta^{\oplus}(q,k,q')$ Acceptance $\forall f \in F, \Delta_f^{\oplus}(q,k,q')$ $\forall f \in F, \Delta_f^{\oplus}(q,k,q')$ Initial states $\forall f \in F, \Delta_f^{\oplus}(q,k,q')$ $\forall f \in F, \Delta_f^{\oplus}(q,k,q')$ Initial states $Transitions$ $S_0(s)$ $F^{\oplus}(s,k,s')$ Initial states $S_0(s)$ $S_0(s)$ $S_0(s)$ Initial states $S_0(s)$ $S_0(s)$ $S_0(s)$ Initial states $S_0(s,q',s',q')$ $S_0(s)$ $S_0(s)$ Initial states $S_0(s,q',s',q')$ $S_0(s)$ $S_0(s)$ Initial states $S_0(s,q',s',q')$ $S_0(s)$ $S_0(s)$ Initial states $S_0(s,q',s')$ $S_0(s)$ $S_0(s)$ Initial states $S_0(s,q',s')$ $S_0(s,d',s')$ $S_0(s)$ Initial states $S_0(s,q',s')$ $S_0(s,d',s')$ $S_0(s,d',s')$ Initial states $S_0(s,q',s')$ $S_0(s,d',s')$ $S_0(s,d',s')$ Initial states $S_0(s,q',s')$ $S_0(s,d',s')$ $S_0(s,d',s')$ Initial states $S_0(s,d',s')$ $S_0(s,d',s')$ $S_0(s,d',s')$ Initial sta	uo	Initial states	$\mathcal{Q}_0(q)$		$U_0(q,\ell)$	
$ \begin{array}{ c c c c } \hline \textbf{A} & \textbf{Acceptance} & \forall f \in F, \Delta_f(q,\ell,q') & \forall f \in F, \Delta_f^\oplus(q,k,q') \\ \hline \textbf{List black etter condimunations} & \forall f \in F, \Delta_f^\oplus(q,k,q') & \forall f \in F, \Delta_f^\oplus(q,k,q') \\ \hline \textbf{List black etter condimunations} & \textbf{R}(s,s') & \textbf{R}(s,s'$	16motu	Transitions	$\Delta(q,\ell,q')$		$\Delta^\oplus(q,k,q')$	
$ \left\{ \begin{array}{c c c c c c c c c c c c c c c c c c c $	Ŋ	Acceptance	$\forall f \in F, \Delta_f(q,\ell,q')$		$\forall f \in F,  \Delta^\oplus_f(q,k,q')$	
State labels $R(s,s')$ $R(s,s')$ $R^{\oplus}(s,k,s')$ Kate labels $R(s,s')$ $R(s,s') \in R \land (L(s) \oplus L(s')) = k$ )Low $L(s,\ell)$ $L(s,\ell)$ Initial states $P_0(s,q) =$ $S_0(s) \land Q_0(q)$ $\exists k, \ell, \ell' [R(s,s') \land \Delta^{\oplus}(q,k,q') \land \Delta^{\oplus}(q,k,s') \land \Delta^{\oplus}(q,k,q')]$ Transitions $T(s,q,s',q') =$ $\exists \ell[R(s,s') \land L(s',\ell) \land \Delta^{\oplus}(q,k,q') \land \Delta^{\oplus}(q,k,s') \land \Delta^{\oplus}(q,k,q')]$ $\exists \ell[R^{\oplus}(s,k,s') \land \Delta^{\oplus}(q,k,q')]$ Spechtance $T_f(s,q,s',q') =$ $\exists \ell[R(s,s') \land L(s',\ell) \land \Delta^{\oplus}(q,k,q') \land \Delta^{\oplus}(q,k,q') \land \Delta^{\oplus}(q,k,q')]$ $\exists k[R^{\oplus}(s,k,s') \land \Delta^{\oplus}(q,k,q')]$ Spechtance $T_f(s,q,s',q') =$ $\exists \ell[R(s,s') \land L(s',\ell) \land \Delta^{\oplus}(q,k,q') \land \Delta^{\oplus}(q,k,q') \land \Delta^{\oplus}(q,k,q')]$ $\exists k[R^{\oplus}(s,k,s') \land \Delta^{\oplus}(q,k,q')]$	-Jou	Initial states			$S_0(s)$	
$ \begin{array}{ c c c c c c } \hline \mathbf{K} & \text{State labels} & & & & & & & & & & & & & & & & & & &$	ıts əxdi	Transitions		R(s,s')	$(\text{encoding } (s,s') \in$	(s,k,s') $R \land (L(s) \oplus L(s')) = k)$
$ \begin{array}{ c c c c c c c c } \hline \mathbf{F} & \mathbf{F}$	Кц	State labels			$L(s,\ell)$	
$ \begin{array}{ c c c c c c c c c c c c c c c c c c c$	1000	Initial states $P_0(s,q) =$	$S_0(s) \wedge \mathcal{Q}_0(q)$		$\exists \ell[S_0(s) \land L(s,\ell) \land U_0(q,\ell)]$	
$ \left  \begin{array}{c} \overrightarrow{P} \\ \overrightarrow{P} \\ \text{Acceptance } T_f(s,q,s',q') = \\ 1 \\ \neg \Delta_f(g,\ell,q') \\ - \Delta_f(g,\ell,q') \\ \end{array} \right  = \left  \begin{array}{c} \exists \ell[R(s,s') \land \Delta_{\oplus}^{\oplus}(q,k,q') \land A_{\oplus}^{\oplus}(q,k,q') \land A_{\oplus}^{\oplus}(q,k,q') \land A_{\oplus}^{\oplus}(q,k,q') \land A_{\oplus}^{\oplus}(q,k,q') \land A_{\oplus}^{\oplus}(q,k,q') \\ - \Delta_f(g,\ell,q') \\ - \Delta_$	ր։։ հւզ	Transitions $T(s,q,s',q') =$	$\exists \ell[R(s,s') \wedge L(s,\ell) \ \wedge \Delta(q,\ell,q')]$	$\exists k, \ell, \ell' \big[ R(s, s') \land \Delta^{\oplus}(q, k, q') \land L(s, \ell) \land L(s', \ell') \land \oplus (\ell, k, \ell') \big]$	$\exists k[R^{\oplus}(s,k,s') \wedge \Delta^{\oplus}(q,k,q')]$	$\frac{[R^{\oplus}(s, \boldsymbol{0}, s') \land equal(q, q')] \lor}{(\exists k [R^{\oplus}_*(s, k, s') \land \Delta^{\oplus}_*(q, k, q')])}$
	ous	Acceptance $T_f(s,q,s',q') =$	$\exists \ell[R(s,s') \wedge L(s,\ell) \ \wedge \Delta_f(q,\ell,q')]$	$ \exists k, \ell, \ell' \left[ R(s, s') \land \Delta_f^\oplus(q, k, q') \land L(s, \ell) \land \oplus(\ell, k, \ell') \right] $	$\exists k[R^{\oplus}(s,k,s])$	$(h) \wedge \Delta^\oplus_f(q,k,q')]$

Table 6.1: Predicates encoding the automata, the Kripke structures and their synchronous product in the different symbolic approaches.

#### 6.4.1 Implementation

We have implemented the symbolic approach based on TGTA, Changeset-based symbolic Kripke structure and Saturation in LTL-ITS<sup>1</sup> tool, which already provides an implementation of the symbolic approach based on TGBA and Saturation. This tool is built on top of three libraries<sup>2</sup>: SDD/ITS, Spot, and LTSmin.

**Spot** is a model-checking library providing several bricks that can be combined to build model checkers [36]. In our implementation, we reused the modules providing a translation from an LTL formula into a TGBA and into a TGTA [6].

**SDD/ITS** is a library for symbolic representation of state-spaces in the form of Instantiable Transition Systems (ITS): an abstract interface for symbolic Labeled Transition Systems (LTS).

The symbolic encoding of ITS is based on Hierarchical Set Decision Diagrams (SDD) [85]. SDDs allow a compact symbolic representation of states and transition relation.

The algorithms presented in this paper can be implemented using any kind of decision diagram (such as OBDD), but use of the SDD software library allows to easily benefit from the automatic saturation mechanism described in [51].

**LTSmin** [12] can generate state spaces from various input formalisms ( $\mu$ CRL, DVE, GNA, MAPLE, PROMELA, ...) and store the obtained LTS in a concise symbolic format, called Extended Table Format (ETF). We used LTSmin to convert DVE models into ETF for our experiments. This approach offers good generality for our tool, since it can process any formalism supported by LTSmin tool.

**Our symbolic model checker** inputs an ETF file and an LTL formula. The LTL formula is converted into TGBA or TGTA which is then encoded using an ITS. The ETF model is also symbolically encoded using an ITS (see Section 6.4.2). The two obtained ITSs are then composed to build a symbolic product, which is also an ITS. Finally, the OWCTY emptiness check is applied to this product.

In all the approaches evaluated in this experimentation, the symbolic products are encoded using the same variables ordering: we used the *top ordering* proposed by Sebastiani et al. [79]. In top ordering, the variables that encode the property automaton (TGBA or TGTA) are at the top of the variable ordering, they precede the variables encoding the model in the decision diagram of the product.

## 6.4.2 Using ETF to build the transition relation of a changeset-based symbolic Kripke structure

An ETF file<sup>3</sup> produced by LTSmin is a text-based serialization of the symbolic representation of the transition relation of a model whose states consist in n integer variables. Transitions are described in the following tabular form:

0/1 0/1 \*

<sup>1</sup>http://ddd.lip6.fr

\*

<sup>&</sup>lt;sup>2</sup>Respectively http://ddd.lip6.fr, http://spot.lip6.fr, and http://fmt.cs.utwente.nl/tools/ ltsmin.

<sup>&</sup>lt;sup>3</sup>http://fmt.cs.utwente.nl/tools/ltsmin/doc/etf.html

#### 1/2 \* 0/1 \*

. . .

where each column correspond to a variable, and each line describes the effect of a symbolic transition on the corresponding variables. The notation "in/out" means that the variable must have the value "in" for the transition to fire, and the value is then updated to "out". A "\*" means that the variable is not consulted or updated by the transition. Each line may consequently encode a set of explicit transitions that differ only by the values of the starred variables: the support of a transition is the set of unstarred variables.

A changeset-based symbolic Kripke structure, as defined in Section 6.3.3, can be easily obtained from such a description. To obtain a changeset associated to a line in the file, it is enough to compute the difference between values of atomic propositions associated to the *in* variables and the values associated to the *out* variables. Because they do not change, starred variables have no influence on the changeset.

Note that an empty changeset does not necessarily correspond to a line where all variables are starred. Even when *in* and *out* values are different, they may have no influence on the atomic propositions, and the resulting changeset may be empty. For instance if the only atomic proposition considered is  $p = (v_1 > 1)$  (where  $v_1$  denotes the first-column variable), then the changeset associated to the first line is  $\emptyset$ , and the changeset for the second line is  $\{p\}$ .

ETF example for the Dining philosophers model: In the dining philosophers problem, n philosophers spend their lives just thinking and easting. In the dining room, there are n chairsa around a table with a big plate of spaghetti and only n forks available. Each philosopher starts by thinking and when he gets hungry, he sits down and try to pick up the two forks that are closest to him. A philosopher can eat only if he pick up both forks. Then, when a philosopher finishes eating, he puts down the forks and returns to think.

The table below shows a part of the ETF file describing the Dining philosophers model with n = 2: In columns, we have the four variables of the state vector. The first line gives the variables values in the initial state and the rest of the table lists the transitions. Each line describe one ETF transition and each cell of this line indicate the change in the value of the column variable: For example in the first ETF transition, the notation "0/1" in the first column means that the first philosopher *phil\_*0 changes from "0:think" to "1:one" (i.e picks up one fork); the "\*" in the last column means that the value of the second philosopher *phil\_*1 does not change in this transition.

	phil_0	fork[0]	fork[1]	phil_1	phil_0 phil_1 values
Initial state	0	0	0	0	0:think
	0/1	0/1	*	*	1:one
Transitions	1/2	*	0/1	*	2:eat
		• •	•		3:finish

This compressed format of transitions allowed us to easily compute the transition relation of changeset-based symbolic Kripke structure (i.e. computing the predicate  $R^{\oplus}(s,k,s')$  introduced in section. 6.3.3). To illustrate the computing of the predicate  $R^{\oplus}(s,k,s')$  (section. 6.3.3) in the Dining philosophers example, we assume that the set of observed atomic propositions  $AP = \{p_0, p_1\}$  where  $p_0 = "phil_0 = eat"$  and  $p_1 = "phil_1 = eat"$ . The changeset of the first ETF transition is equal to  $\emptyset$  because *phil\_0* does not take the value "2:eat" (i.e.  $p_0$  is false) before and after the

transition,  $p_1$  also does not change because the value of *phil\_1* does not change in this transition. Similarly, it is easy to compute the changeset of the second ETF transition equal to  $\{p_0\}$ .

6.4.3 Benchmar	K
----------------	---

BEEM model	$\frac{\text{states}}{10^3 \times}$	stut. ratio	BEEM model	$\frac{\text{states}}{10^3 \times}$	stut. ratio
at.5	31 999	96%	lann.6	144 151	44%
bakery.4	157	76%	lann.7	160 025	52%
bakery.7	29 047	80%	lifts.7	5 1 2 6	85%
bopdp.3	1 040	86%	lifts.8	12359	85%
brp2.3	40	73%	mcs.5	60 5 56	81%
elevator.4	888	63%	peterson.5	131 064	73%
fischer.5	101 028	82%	pgm_protocol.8	3 0 6 9	88%
iprotocol.7	59794	81%	phils.7	71934	89%
lamport_nonatomic.5	95 118	86%	production_cell.6	14 520	78%
lamport.7	38717	91%	reader_writer.3	604	81%

Table 6.2: Characteristics of our selected benchmark models. The stuttering-ratio represents the percentage of stuttering transitions in the model. Since the definition of stuttering depends on the atomic propositions of the formula, we give an average over the 200 properties checked against each model.

We evaluated the TGBA and TGTA approaches on the following models and formulas:

- Our models come from the BEEM benchmark [67], a suite of models for explicit model checking, which contains some models that are considered difficult for symbolic model checkers [12]. Table 6.2 summarizes the 20 models we used to evaluate our approaches. These models represent various controllers, communication protocols, mutual exclusion and leader election algorithms.
- BEEM provides a few LTL formulas, but they mostly represent safety properties and can thus be checked without building a product. Therefore, for each model, we randomly generated 200 stutter-invariant LTL formulas: 100 verified formulas (empty product) and 100 violated formulas (non-empty product). We consequently have a total 4000 pairs of (model, formula).

All tests were run on a 64bit Linux system running on an Intel(R) 64-bit Xeon(R) @2.00GHz with 64GB of RAM. Executions that exceeded 60 minutes or 6GB of RAM were aborted and are reported with time and memory above these thresholds in our graphics.

#### 6.4.4 Results

The results of our experimental comparisons are presented by the two scatter plot matrices of Figure 6.2 and Figure 6.3.

Figure 6.2 compares the time-performance of the TGTA-approach against the reference TGBA approach. In order to show the influence of the saturation technique, we also ran the TGBA



Figure 6.2: Time-comparison of the TGBA and TGTA approaches, with saturation enabled "(sat)" or disabled "(nosat)", on a set of 4000 pairs (*model*, *formula*). Timeouts and Out-of-memory errors are plotted on separate lines on the top or right edges of the scatter plots. Each plot also displays the number of cases that are above or below the main diagonal (including timeouts and out-of-memory errors), i.e., the number of (*model*, *formula*) for which one approach was better than the other. Additional diagonals show the location of  $\times 10$  and /10 ratios. Points are plotted with transparency to better highlight dense areas, and lessen the importance of outliers.



Figure 6.3: Comparison of the memory-consumption of the TGBA and TGTA approaches, with or without saturation, on the same set of problems.



Figure 6.4: Comparison of TGBA and TGTA approaches, without saturation "(nosat)". Timeouts and Out-of-memory errors are plotted on the top or right edges of the scatter plots.



Figure 6.5: Comparison of TGBA and TGTA approaches, with saturation enabled "(sat)". Timeouts and Out-of-memory errors are plotted on the top or right edges of the scatter plots.

			time	(s)			memor	y (MB)	
		TGI	ЗA	TG	ΓA	TGE	A	TGT	Ά
	cases	(nosat)	(sat)	(nosat)	(sat)	(nosat)	(sat)	(nosat)	(Si
at.5	200	391.86	446.56	590.44	182.95	1668	1649	1760	13
bakery.4	200	14.07	14.74	17.90	7.43	331	341	388	<u> </u>
bakery.7	192	634.79	632.04	593.24	168.80	1353	1361	1330	1
bopdp.3	200	18.89	19.36	28.65	9.44	325	332	392	2
brp2.3	200	5.05	5.20	6.31	2.65	147	149	170	
elevator.4	199	154.02	169.32	204.15	81.40	1142	1144	1165	9
fischer.5	190	300.82	372.75	418.96	187.41	1455	1486	1496	15
iprotocol.7	145	1790.41	1776.72	1987.39	290.30	2271	2302	2148	12
lamport_nonatomic.5	200	111.24	147.24	191.97	17.00	917	935	1108	ധ
lamport.7	196	479.56	570.71	696.13	63.61	1523	1547	1620	$\sim$
lann.6	129	504.95	576.05	793.42	148.53	1073	1048	1021	6
lann.7	196	161.05	208.84	379.38	99.07	944	960	1027	9
lifts.7	200	541.09	579.36	643.38	152.58	1302	1318	1349	9
lifts.8	173	1804.29	1858.65	1910.90	283.42	1903	1891	1919	$\Xi$
mcs.5	187	1659.96	1889.17	2208.65	460.33	4155	3890	3718	20
peterson.5	199	196.12	283.39	309.68	55.40	1105	1156	1145	~1
pgm_protocol.8	194	935.45	982.70	1118.97	658.58	1419	1423	1437	13
phils.7	200	6.79	12.33	15.58	1.64	181	200	330	
production_cell.6	197	663.13	702.71	862.59	375.44	1367	1388	1515	14
reader_writer.3	200	0.67	0.71	1.12	0.63	52	53	63	
all models	3797	487.65	530.94	612.79	158 04	1204	1201	1232	ø

and memory consumption is computed over the 3797 cases where all methods terminated normally (without timeout or out-of-memory error). Table 6.3: Comparison of the performances of the symbolic approaches TGBA versus TGTA , presented model by model. The average run-time and TGTA approaches with saturation disabled. In our comparison matrix, the labels "(sat)" and "(nosat)" indicate whether saturation was enabled or not. Each point of the scatter plots represents a measurement for a pair (*model*, *formula*). Axes use a logarithmic scale. The colors distinguish violated formulas (non-empty products) from verified formulas (empty products).

In each scatter plot, each point below the diagonal is in favor of the approach displayed on the right, while each point above the diagonal is in favor of the approach displayed in the top of the scatter plot. For instance, the highlighted scatter plot (bottom of Figure 6.2) compares the time-performance between TGBA and TGTA approaches where the saturation is enabled (sat), the x-axis represents the "TGBA (sat)" approach and the y-axis represents the "TGTA (sat)" approach, so 3424 points below the diagonal correspond to cases where the "TGTA (sat)" approach is better, and the 548 points above the diagonal corresponds to points were the "TGBA (sat)" approach is better.

Figure 6.3 gives the memory view of this experiment.

As shown by the highlighted scatter plots in Figure 6.2 and 6.3, the TGTA approach clearly outperforms the traditional TGBA-based scenario by one order of magnitude. This is due to the combination of two factors: saturation and exploration of stuttering.

The saturation technique does not significantly improve the model checking using TGBA (compare "TGBA (sat)" against "TGBA (nosat)" at the top of Figure 6.2 and 6.3). In fact, the saturation technique is limited on the TGBA approach, because in the transition relation of Definition 34 each conjunction must consult the variable q representing the state of the TGBA, therefore q impacts the supports and the reordering of clusters evaluated by the saturation. This situation is different in the case of TGTA approach, where the  $T_{\emptyset}$  term of the transition-relation of the product (equation (6.4)) does not involve the state q of the TGTA: here, saturation strongly improves performances (compare "TGTA (sat)" against "TGTA (nosat)").

Overall the improvement to this symbolic technique was only made possible because the TGTA representation makes it easy to process the stuttering behaviors separately from the rest. These stuttering transitions represent a large part of the models transitions, as shown by the stuttering-ratios of Table 6.2. Using these stuttering-ratios, we can estimate in our Benchmark the importance of the term  $T_0$  compared to  $T_*$  in equation (6.4).

Figures 6.4 and 6.5 show more clearly the most important scatter plots of the two matrices. Figure 6.4 shows the scatter plots (for time and memory performance) comparing TGBA and TGTA approaches with saturation disabled (nosat). Figure 6.5 displays the same comparison but with saturation enabled (sat).

Table 6.3 gives an overview of the performance of the TGBA and TGTA approach, model by model. The average run-time and memory consumption is computed over the 3797 cases where all methods terminated normally (without timeout or out-of-memory error).

Table 6.4 shows the best and the worst approach among the four possible combinations, i.e., TGBA and TGTA approaches with saturation disabled (no saturation) or with saturation enabled. The verified formulas are separated from violated formulas.

The Tables 6.3 and 6.4 confirm the above observations deduced from the scatter plots: i.e., the TGTA approach with saturation outperform the TGBA approach (with or without saturation).

		no saturation		saturation	
		TGBA	TGTA	TGBA	TGTA
Verified	Win	297 (14%)	13 (0%)	222 (11%)	1613 (81%)
	Lose	204 (10%)	1359 (68%)	464 (23%)	67 (3%)
	Fail	29 (1%)	43 (2%)	30 (1%)	2 (0%)
Violated	Win	150 (7%)	1 (0%)	97 (4%)	1795 (89%)
	Lose	118 (5%)	1594 (79%)	297 (14%)	26 (1%)
	Fail	12 (0%)	89 (4%)	32 (1%)	7 (0%)

Table 6.4: On all experiments (grouped by verified versus violated formulas), we count the number of cases a specific method has (Win) the best time or (Lose) it has either run out of time/memory or it has the worst time amongst successful methods. The Fail line shows how much of the Lost cases failed (because of timeout or out-of-memory). The sum of a line may exceed 100% if several methods are equally placed.

## 6.5 Conclusion

Testing automata are a way to improve the explicit model checking approach when verifying stutter-invariant properties, but they had not been used for symbolic model checking.

In the previous chapters, we introduced a generalization of these testing automata, called TGTA, and we evaluated their use for explicit model checking.

In this chapter, we have shown how to use TGTA in a symbolic approach. We compare this new TGTA approach to a more classical symbolic approach (using TGBA).

On our benchmark, using TGTA and saturation technique, we were able to gain one order of magnitude over the TGBA-based approach.

Three versions of this TGTA approach, including a basic version and two improvements were implemented. The first improvement was to use a new encoding of a Kripke structure inspired by the TGTA transition relation, i.e., based on the labeling of transitions by changesets. The second improvement is based on the exploration of stuttering transitions during the emptiness check of the symbolic product.

The latter optimization is based on the property of TGTA that all stuttering transitions are self-loops and every state has a stuttering self-loop. Consequently, the exploration of stuttering transitions in the product is equivalent to explore stuttering transitions in the model (remaining in the same TGTA states).

Using this property, we have shown that fixpoints over the transition relation of a product between a Kripke structure and a TGTA can benefit from the saturation technique, especially because part of their expression is only dependent on the model, and can be evaluated without consulting the transition relation of the property automaton. This allows to the saturation algorithm to ignore the symbolic variables encoding the TGTA in the product, and therefore effectively saturate the symbolic product nodes of the variables encoding the model.

This improvement was possible only because TGTA makes it possible to process stuttering behaviors specifically, in a way that helps the saturation technique.

In the next chapter, we evaluate the use of TGTA in the context of hybrid approaches, which

combines the use of both explicit and symbolic approaches [79, 38]. These hybrid approaches use an On-the-Fly exploration unlike the symbolic approaches presented in this chapter.

## CHAPTER 7

# Hybrid LTL Model Checking using TGTA

#### Contents

7.1	Introduction						
7.2	7.2 Preliminaries						
	7.2.1	TGBA labeled with propositional formulas 135					
	7.2.2	TGTA labeled with propositional formulas 136					
7.3	7.3 Symbolic Observation Graph (SOG)						
	7.3.1	SOG 141					
	7.3.2	SOG for TGTA (SOG-TGTA) 142					
7.4	7.4 Symbolic Observation Product (SOP)						
	7.4.1	SOP					
	7.4.2	SOP Using TGTA (SOP-TGTA)					
7.5	5 Self-Loop Aggregation Product (SLAP)						
	7.5.1	SLAP					
	7.5.2	SLAP Using TGTA (SLAP-TGTA) 147					
7.6	7.6 Experimental Comparison of Hybrid Approaches using TGBA vs. To						
	7.6.1	Implementation					
	7.6.2	Results					
	7.6.3	SOG versus SOG-TGTA					
	7.6.4	SOP versus SOP-TGTA 151					
	7.6.5	SLAP versus SLAP-TGTA					
7.7	Conclusion						

## 7.1 Introduction

In the previous chapter, we have shown that the symbolic approach to model checking allows to encode the product automaton in a concise way, but its emptiness check can not be performed on-the-fly as in the explicit approach.

In order to take advantage of the best of the both worlds, *Hybrid approaches* [79, 38] to model checking are proposed as combinations of explicit and symbolic approaches.

In these hybrid approaches, the property automaton is described explicitly because its size is not large in most cases (and can be reduced by means of several explicit optimizations). However, the state-space of the model is typically very large and must be encoded symbolically. An hybrid approach is generally based on an on-the-fly construction of an explicit graph of symbolic nodes, called aggregates. Each aggregate symbolically encodes a set of states of the Kripke structure or of the product. In this work, we focused on three hybrid techniques [37]: SOG, SOP and SLAP. The Symbolic Observation Graph (SOG) approach exploits the fact that only a subset of the atomic propositions of the model are observed by the LTL property to check. A SOG is an abstraction of a Kripke structure where consecutive states are aggregated if they share the same values of the observed atomic propositions. The *Symbolic Observation Product* (SOP) approach tries to allow further aggregation than SOG by exploiting the fact that the number of observed atomic propositions decreases as we progress in the property automaton. A *Self-Loop Aggregation Product* (SLAP) is similar to a SOP, an aggregation graph alternative to the traditional product automaton. In SLAP, the Kripke structure states are aggregated according to the valuations of the self-loops in the property automaton.

As for symbolic model checking, Testing Automata (and their variants) have never been used before for hybrid model checking.

The goal of this chapter is to show how the traditional hybrid approaches based on TGBA can be adapted to obtain TGTA-based approaches.

In this chapter, we present three existing hybrid approaches SOG, SOP and SLAP that are based on TGBA. Then, we define and implement variations of these three approaches using TGTA instead of TGBA. We then experimentally compare the performance of each hybrid approach (SOG, SOP and SLAP) against its TGTA-based variant (SOG-TGTA, SOP-TGTA and SLAP-TGTA).

## 7.2 Preliminaries

The formalization of hybrid approaches requires the following definitions and notations introduced by Duret-Lutz et al. [38].

We firstly remind the formalization of propositional formula, then we present alternative definitions of TGBA Duret-Lutz et al. [38] and TGTA that use these propositional formulas as labels.

- $\mathbb{B} = \{\bot, \top\}$  represents the Boolean values.
- B(*AP*) is the set of all propositional formulas over *AP*. The formulas of B(*AP*) are built inductively from the propositions *AP*, B, and the logical operators ∧, ∨, and ¬.
- Using this definition of  $\mathbb{B}(AP)$ ,  $AP' \subseteq AP$  implies that  $\mathbb{B}(AP') \subseteq \mathbb{B}(AP)$ .
- *FV*(φ) (Free Variables) Duret-Lutz et al. [38] is the set of atomic propositions observed in the formula φ, e.g., *FV*(*ab* ∨ *ab*) = {*a*} because φ = (*ab* ∨ *ab*) can be simplified into φ = *a* (we say that *b* is a silent (or a "don't care") atomic proposition in φ).

• The notation  $\ell \stackrel{AP'}{=} \ell'$  is equivalent to  $\ell_{|AP'} = \ell'_{|AP'}$ , where  $\ell_{|AP'}$  denotes the restriction of the valuation  $\ell$  to the subset of atomic propositions AP'. In other words,  $\ell \stackrel{AP'}{=} \ell'$  means that the valuations  $\ell$  and  $\ell'$  match on the atomic propositions of AP'.

Duret-Lutz et al. [38] use an alternative definition of TGBA more suited than Definition 16 to define and implement the hybrid approaches SOP and SLAP. In these TGBA, each transition is labeled with a propositional formula  $\phi$  over *AP* (the resulting transition relation is of the form  $\delta \subseteq Q \times \mathbb{B}(AP) \times 2^{\mathcal{F}} \times Q$ ). In the following, we will see that this labeling of transitions with propositional formula simplifies the formalization of the different hybrid approaches presented in this work. For instance, in SOP and SOP-TGTA approaches, it allows to simplify the definition (and the implementation) of the concept of observed alphabet from a state of an automaton (TGBA or TGTA).

#### 7.2.1 TGBA labeled with propositional formulas

In a TGBA labeled with propositional formulas, each transition is labeled with a propositional formula  $\phi \in \mathbb{B}(AP)$  instead of a valuation  $\ell \in 2^{AP}$ . This formulas  $\phi$  represents the maximal set of valuations  $\{\ell_0, \ell_1, \dots, \ell_n\} \subseteq 2^{AP}$  such that  $\forall i \leq n, \ell_i \models \phi$ . In other words, the set of valuations  $\{\ell_0, \ell_1, \dots, \ell_n\}$  is the set of all models of  $\phi$  (it can be represented by  $\phi = \bigvee_{i \leq n} \ell_i$ )). For example, in Figure 7.1a, the formula  $\phi = b$  labeling the transition  $q_0 \xrightarrow{b} q_1$  represents the set of valuations  $\{\bar{a}b, ab\}$ .

In order to obtain this form of TGBA, we merge into a single transition all transitions between each pair of states (q,q'), when these transitions are labelled with the same acceptance conditions F. Formally, for a pair of states  $(q,q') \in Q^2$  and a set of acceptance conditions  $F \in \mathcal{F}$ , the set of transitions  $(q, \ell_0, F, q'), \ldots, (q, \ell_n, F, q') \in \delta$  are merged into a single transition  $(q, \phi, F, q')$  where  $\phi = \bigvee_{i \le n} \ell_i$ .

**Definition 39** (TGBA (labeled with propositional formulas)). A Transition-based Generalized Büchi Automata (TGBA) over the alphabet  $\Sigma = 2^{AP}$  is a tuple  $G = \langle Q, I, \delta, \mathcal{F} \rangle$  where

- *Q* is a finite set of states,
- *I is a finite set of initial states,*
- $\mathcal{F} \neq \emptyset$  is a finite and non-empty set of acceptance conditions,
- $\delta \subseteq Q \times \mathbb{B}(AP) \times 2^{\mathcal{F}} \times Q$  is a transition relation, where each element  $(q, \phi, F, q') \in \delta$  represents a transition from state q to state q' labeled by a propositional formula  $\phi \in \mathbb{B}(AP)$ , and a set of acceptance conditions  $F \in 2^{\mathcal{F}}$ . In the following, an element  $(q, \phi, F, q') \in \delta$  will be denoted  $q \xrightarrow{\phi, F} q'$

An infinite word  $\sigma = \ell_0 \ell_1 \ell_2 \ldots \in \Sigma^{\omega}$  is accepted by  $\mathcal{G}$  if there exists an infinite sequence of transitions  $\pi = (q_0, \phi_0, F_0, q_1)(q_1, \phi_1, F_1, q_2) \cdots (q_i, \phi_i, F_i, q_{i+1}) \cdots \in \delta^{\omega}$  ( $\pi$  is called a run of  $\mathcal{G}$ ) where:

•  $q_0 \in I$ , and  $\forall i \in \mathbb{N}, \ell_i \models \phi_i$  (i.e., the infinite word  $\sigma$  is recognized by the run  $\pi$ )

∀f ∈ 𝔅, ∀i ∈ ℕ, ∃j ≥ i, f ∈ F<sub>j</sub> (i.e., the run π is accepting iff it visits each acceptance condition infinitely often).

*The language accepted by* G *is the set*  $\mathscr{L}(G) \subseteq \Sigma^{\omega}$  *of infinite words it accepts.* 

Figure 7.1a shows the TGBA of the LTL formula  $a \cup b$ . The transitions are labeled with the propositional formulas  $a\bar{b}$ , b and  $\top$  (which encodes all the valuations over  $AP = \{a, b\}$ ).

Using this variant of TGBA, we obtain a new definition of the synchronous product between a TGBA  $\mathcal{G}$  and a Kripke structure  $\mathcal{K}$ . We remind that this synchronous product is also a TGBA, which only accepts the words accepted by both  $\mathcal{G}$  and  $\mathcal{K}$  (Formally,  $\mathscr{L}(\mathcal{K} \otimes \mathcal{G}) = \mathscr{L}(\mathcal{K}) \cap \mathscr{L}(\mathcal{G})$ ).

**Definition 40** (Synchronous product of a TGBA with a Kripke structure). For a TGBA  $\mathcal{G} = \langle Q, I, \delta, \mathcal{F} \rangle$  over the alphabet  $\Sigma = 2^{AP}$  and a Kripke structure  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$ , the product  $\mathcal{G} \otimes \mathcal{K}$  is the TGBA  $\langle Q_{\otimes}, I_{\otimes}, \delta_{\otimes}, \mathcal{F} \rangle$  over the alphabet  $\Sigma = 2^{AP}$  where

- $Q_{\otimes} = Q \times S$ ,
- $I_{\otimes} = I \times \mathcal{S}_0$ ,
- $\delta_{\otimes} \subseteq Q_{\otimes} \times \mathbb{B}(AP) \times 2^{\mathcal{F}} \times Q_{\otimes}$  where

$$\boldsymbol{\delta}_{\otimes} = \left\{ (q_1, s_1) \xrightarrow{l(s_1), F} (q_2, s_2) \middle| \begin{array}{c} (s_1, s_2) \in \mathcal{R} \text{ and} \\ \exists \boldsymbol{\phi} \in \mathbb{B}(AP) \text{ s.t. } q_1 \xrightarrow{\boldsymbol{\phi}, F} q_2 \in \boldsymbol{\delta} \text{ and } l(s_1) \models \boldsymbol{\phi} \end{array} \right\}$$

Figure 7.1c (taken from [37]) is an illustration of Definition 40. It shows an example of a synchronous product  $\mathcal{G} \otimes \mathcal{K}$  between a TGBA  $\mathcal{G}$  of  $a \cup b$  (Figure 7.1a) and an example of Kripke structure  $\mathcal{K}$  over  $AP' = \{a, b, c\}$  shown in Figure 7.1b. The initial state of the product is  $(q_0, s_0)$ . Then, the successors  $\{s_1, s_4\}$  of  $s_0$  in  $\mathcal{K}$  are synchronized with the state  $q_0$  of  $\mathcal{G}$ , because the TGBA self-loop  $q_0 \xrightarrow{d\bar{b}} q_0$  is labeled by the formula  $\phi = a\bar{b}$  and  $l(s_0) = a\bar{b}c \models \phi$ . From state  $(q_0, s_4)$ , the product move to state  $(q_1, s_5)$  through the TGBA transition  $q_0 \xrightarrow{b} q_1$  because  $l(s_4) = ab\bar{c} \models b$ . From the product state  $(q_1, s_5)$ , the TGBA state  $q_1$  only requires to verify  $\top$  (i.e, any valuation) to explore the self-loop labeled with the acceptance condition  $\bullet$ . Therefore, any cycle of  $\mathcal{K}$  starting in  $s_5$  corresponds to an accepting cycle in the product.

#### 7.2.2 TGTA labeled with propositional formulas

Similar to TGBA, we present in this section a definition of a TGTA labeled with propositional formulas, and for the same reasons as for TGBA, we will see that this alternative definition of TGTA is more suited than Definition 30 to define and implement the TGTA-based hybrid approaches presented later in this chapter.

The difference between TGBA and TGTA is mainly in the interpretation of the transition relation  $\delta \subseteq Q \times \mathbb{B}(AP) \times 2^{\mathcal{F}} \times Q$ . Indeed, in TGTA labeled with propositional formulas, for each transition  $(q, \phi, F, q') \in \delta$ , the formula  $\phi$  encodes a set of changesets  $\{k_0, k_1, \ldots, k_n\} \subseteq 2^{AP}$ , where  $\forall i \leq n$ , each changeset  $k_i \models \phi$ .

As for TGBA, in order to obtain the TGTA labeled with propositional formulas, we merge into a single transition all transitions between each pair of states (q, q'), when these transitions are


abīc

аbc

Figure 7.1: Examples [37] Using TGBA

аbc

 $ab\bar{c}$ 



(a) TGTA  $\mathcal{T}$  for  $a \cup b$ . The formula  $\bar{a}\bar{b}$  labeling the self-loop on  $q_0$  represents the empty changeset  $\emptyset$ . The formula *b* labeling the transition  $q_0 \xrightarrow{b} q_1$ represents the set of changesets  $\{\{b\}, \{a, b\}\}$ . The label  $\top$  of the self-loop on  $q_1$  encodes the set of all changesets  $\{\emptyset, \{a\}, \{b\}, \{a, b\}\}$ .



(b) Kripke structure  $\mathcal{K}$ 



(c) 0-TGTA of product  $\mathcal{T}\otimes\mathcal{K}$ 





(d) Symbolic Observation Graph for TGTA: a SOG-TGTA  $\widehat{\mathcal{K}}'_{\{a,b\}}$ 







(g) 0-TGTA of the SLAP  $\mathcal{T}\boxtimes \mathcal{K}$ 

Figure 7.2: Examples Using TGTA

labelled with the same acceptance conditions F. Formally, , for a pair of states  $(q,q') \in Q^2$  and a set of acceptance conditions  $F \in \mathcal{F}$ , the set of transitions  $(q,k_0,F,q'),\ldots,(q,k_n,F,q') \in \delta$  are merged into a single transition  $(q,\phi,F,q')$  where the set of changesets  $\{k_0,\ldots,k_n\}$  is the set of all models of  $\phi$ . In the following, we will use  $(q,\phi,F,q')$  and  $(q,\{k_0,\ldots,k_n\},F,q')$  interchangeably as transitions of  $\delta$ . For example, in Figure 7.2a, the formula  $\phi = b$  labeling the transition  $q_0 \xrightarrow{b} q_1$ represents the set of changesets  $\{\{b\},\{a,b\}\}$  (meaning that the value of b changes between  $q_0$  and  $q_1$  and we "do not care" about a).

The following definition formalizes this form of TGTA and how it changes the way an infinite word is accepted.

**Definition 41** (TGTA (labeled with propositional formulas)). *A Transition-based Generalized Testing Automata (TGTA) over the alphabet*  $\Sigma = 2^{AP}$  *is a tuple*  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F} \rangle$  *where* 

- *Q* is a finite set of states,
- I is a finite set of initial states,
- U: I → B(AP) is a function mapping each initial state to a propositional formula \$\$\$\$ € B(AP),
- $\mathcal{F} \neq \emptyset$  is a finite and non-empty set of acceptance conditions,
- $\delta \subseteq Q \times \mathbb{B}(AP) \times 2^{\mathcal{F}} \times Q$  is a transition relation, where each element  $(q, \phi, F, q') \in \delta$  represents a transition from state q to state q' labeled by a propositional formula  $\phi \in \mathbb{B}(AP)$ , and a set of acceptance conditions  $F \in 2^{\mathcal{F}}$ .
- $\delta$  has to satisfy the following stuttering-normalization constraint:  $\forall (q,q') \in Q_{*}^{2} : (\exists (\phi,F) \in \mathbb{B}(AP) \times 2^{\mathcal{F}}, \emptyset \models \phi \land (q,\phi,F,q') \in \delta) \iff (q=q')$

An infinite word  $\sigma = \ell_0 \ell_1 \ell_2 \ldots \in \Sigma^{\omega}$  is accepted by  $\mathcal{T}$  iff there exists an infinite sequence of transitions  $\pi = (q_0, \phi_0, F_0, q_1)(q_1, \phi_1, F_1, q_2) \cdots (q_i, \phi_i, F_i, q_{i+1}) \cdots \in \delta^{\omega}$  ( $\pi$  is called a run of  $\mathcal{T}$ ) where:

- $q_0 \in I$ , and  $\ell_0 \models U(q_0)$
- $\forall i \in \mathbb{N}, (\ell_i \oplus \ell_{i+1}) \models \phi_i$  (*i.e.*, the infinite word  $\sigma$  is recognized by the run  $\pi$ )
- $\forall f \in \mathcal{F}, \forall i \in \mathbb{N}, \exists j \ge i, f \in F_i \text{ (i.e., the run } \pi \text{ is accepting).}$

*The language accepted by* T *is the set*  $\mathscr{L}(T) \subseteq \Sigma^{\omega}$  *of infinite words it accepts.* 

Figure 7.2a shows the TGTA of the LTL formula  $a \cup b$ . The valuations of the initial states are  $U(q_0) = a\bar{b}$  and  $U(q_1) = b = \{ab, \bar{a}b\}$ . The transitions are labeled with the propositional formulas. The formula  $\bar{a}\bar{b}$  labeling the stuttering self-loop on  $q_0$  is the empty changeset  $\emptyset$ . The formula b labeling the transition  $q_0 \xrightarrow{b} q_1$  is obtained by merging the two changesets of the transitions  $q_0 \xrightarrow{\{b\}} q_1$  and  $q_0 \xrightarrow{\{a,b\}} q_1$  (because  $\{\{b\}, \{a,b\}\} \models b$ ). On the self-loop  $q_1 \xrightarrow{\top, \{\bullet\}} q_1$ , the formula  $\top$  is obtained by merging the set of all changesets over  $AP = \{a,b\}$ , and the acceptance condition is indicated by the black dot ( $\mathcal{F} = \{\bullet\}$ ).

Using this alternative definition of TGTA, we obtain the following definition of the synchronous product between a TGTA and a Kripke structure. We remind that this synchronous product is an  $\emptyset$ -TGTA (Definition 29). **Definition 42** (Synchronous product of a TGTA with a Kripke structure). For a TGTA  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F} \rangle$  over the alphabet  $\Sigma = 2^{AP}$  and a Kripke structure  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$ , the product  $\mathcal{T} \otimes \mathcal{K}$  is the 0-TGTA  $\langle Q_{\otimes}, I_{\otimes}, U_{\otimes}, \delta_{\otimes}, \mathcal{F} \rangle$  over the alphabet  $\Sigma = 2^{AP}$  where

- $Q_{\otimes} = Q \times S$ ,
- $I_{\otimes} = \{(q,s) \in I \times \mathcal{S}_0 \mid l(s) \models U(q)\}$
- $\forall (q,s) \in I_{\otimes}, U_{\otimes}((q,s)) = l(s),$
- $\delta_{\otimes} \subseteq Q_{\otimes} \times \mathbb{B}(AP) \times 2^{\mathcal{F}} \times Q_{\otimes}$  where

$$\delta_{\otimes} = \left\{ (q_1, s_1) \xrightarrow{(l(s_1) \oplus l(s_2))_{|AP}, F} (q_2, s_2) \middle| \begin{array}{l} (s_1, s_2) \in \mathcal{R} \text{ and} \\ \exists \phi \in \mathbb{B}(AP) \text{ s.t. } q_1 \xrightarrow{\phi, F} q_2 \in \delta \text{ and} \\ (l(s_1) \oplus l(s_2)) \models \phi \end{array} \right\}$$

Figure 7.2c is an illustration of Definition 42. It shows an example of a synchronous product  $\mathcal{T} \otimes \mathcal{K}$  between the TGTA  $\mathcal{T}$  of  $a \cup b$  of Figure 7.2a and the Kripke structure  $\mathcal{K}$  of Figure 7.1b. The initial state of the product is  $(q_0, s_0)$  because  $l(s_0) = a\bar{b}c \models U(q_0) = a\bar{b}$ . Then,  $(q_0, s_0)$  have two successors: the first successor is  $(q_0, s_1)$  because  $\mathcal{K}$  has a transition  $s_0 \to s_1$  with  $l(s_0) \oplus l(s_1) = \emptyset \models \bar{a}\bar{b}$  and  $\mathcal{T}$  have a stuttering self-loop  $q_0 \stackrel{\bar{a}\bar{b}}{\longrightarrow} q_0$ ; the second successor is  $(q_1, s_4)$  because in  $\mathcal{K}$  we have  $s_0 \to s_4$  with  $l(s_0) \oplus l(s_4) = \{b, c\} \models b$  and the TGTA have a transition  $q_0 \stackrel{b}{\to} q_1$  labeled with  $\phi = b$ . From the product state  $(q_1, s_4)$ , the TGTA can explore any changeset through the self-loop labeled with  $\top$  and the acceptance condition  $\bullet$ . Therefore, the TGTA state  $q_1$  can be synchronized with any reachable state from  $s_4$  in  $\mathcal{K}$  and the cycle  $(q_1, s_4) \to (q_1, s_5) \to (q_1, s_6) \to (q_1, s_7) \to (q_1, s_4)$  is an accepting cycle in the product. In the obtained product  $\mathcal{T} \otimes \mathcal{K}$ , each transition is labeled with the changeset  $((l(s_1) \oplus l(s_2))|_{AP})$  between the states of  $\mathcal{K}$ . These changesets are computed according to the set of atomic propositions  $AP = \{a, b\}$  observed by  $\mathcal{T}$ . The product transitions also bear the acceptance conditions coming from the TGTA  $\mathcal{T}$ .

## 7.3 Symbolic Observation Graph (SOG)

In this section, we propose an adaptation of the SOG hybrid approach for use with TGTA instead of TGBA. We start by presenting a variant of SOG [58, 37] used in a TGBA-based approach [37], then we present another variant of SOG proposed in Klai and Poitrenaud [58], this variant is called SOG-TGTA in this work because it is more suited to be used in a TGTA-based approach. In Section 7.6.3, we will show the results of an experimental comparison between the original approach based on SOG and TGBA and our approach based on TGTA and SOG-TGTA.

A SOG is a transformation of a Kripke structure allowing to aggregate states according to the set AP of atomic propositions observed in the property automaton. This transformation only preserves stutter-invariant properties. The constructed SOG is an explicit graph where each node is a symbolic set of states. Theses states are aggregated because they share the same values for the atomic propositions of AP (they may have different values for the other atomic propositions of the model that are not in AP).

In hybrid approaches, symbolic data structure are used to represent sets of states of the Kripke structure. The following symbolic operations are introduced in [38] to manipulate this symbolic aggregate of states.

Let  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$  a Kripke structure, encoded by a Symbolic Kripke structure  $K = \langle S_0, R, L \rangle$  (Definition 32). For a set of states  $a \subseteq S$  and a propositional formula  $\phi \in \mathbb{B}(AP)$ , the symbolic operations SuccF $(a, \phi)$  and ReachF $(a, \phi)$  are defined as follows:

- SuccF(a, φ) = {s' ∈ S | ∃s ∈ a, R(s,s') ∧ ∃ℓ, [ℓ ⊨ φ ∧ L(s', ℓ)]}, i.e., the set of the Successors of states of aggregate a, Filtered to keep only those satisfying φ.
- ReachF(*a*, φ) computes the least subset of *S* satisfying:
  - $a \subseteq \operatorname{ReachF}(a, \phi)$ ,
  - SuccF(ReachF( $a, \phi$ ),  $\phi$ )  $\subseteq$  ReachF( $a, \phi$ ).

Reach $F(a, \phi)$  is implemented using symbolic least fixed-points on Decision Diagrams.

**Definition 43** (Homogeneous aggregate [37]). Let  $a \in 2^{S} \setminus \{\emptyset\}$  be a subset of states of K. We say that a is a homogeneous aggregate w.r.t. (with respect to) a given set of atomic propositions AP iff  $\forall s, s' \in a, l(s) \stackrel{AP}{=} l(s')$ . In other words, all the states of the aggregate a have the same valuation for all the atomic propositions in AP.

For a homogeneous aggregate a w.r.t. AP, we write  $l_{AP}(a) = l(s)_{|AP|}$  for any state  $s \in a$  (i.e., the valuation of a is the valuation of any one of its states).

For any  $AP' \subseteq AP$ , a homogeneous aggregate a w.r.t. AP is also homogeneous w.r.t. AP'.

#### 7.3.1 SOG

**Definition 44** (Symbolic Observation Graph [58, 37]). Let  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$  be a Kripke structure over the set of atomic propositions  $AP_{\mathcal{K}}$ . A symbolic observation graph over  $AP \subseteq AP_{\mathcal{K}}$  is defined as  $\widehat{\mathcal{K}}_{AP} = \langle S', S'_0, \mathcal{R}', l' \rangle$  a Kripke structure over AP satisfying :

as  $\widehat{\mathcal{K}}_{AP} = \langle \mathcal{S}', \mathcal{S}'_0, \mathcal{R}', l' \rangle$  a Kripke structure over AP satisfying : 1.  $\mathcal{S}' = \Gamma' \cup 2^{AP}$  with  $\Gamma' = \left\{ a \in 2^{\mathcal{S}} \setminus \{\emptyset\} \middle| \begin{array}{l} a \text{ is homogeneous w.r.t. } AP \\ a = \operatorname{ReachF}(a, l_{AP}(a)) \end{array} \right\}$ 

Elements of  $\Gamma'$  are called aggregates and elements of  $2^{AP}$  are divergent states.

2. 
$$\forall a \in \mathcal{S}', \ l'(a) = \begin{cases} l_{AP}(a) & \text{if } a \in \Gamma' \\ a & \text{if } a \in 2^{AP} \end{cases}$$

3. 
$$\mathcal{R}' = \{a \to a' \in \Gamma' \times \Gamma' \mid a' = \operatorname{ReachF}(\operatorname{SuccF}(a, l'(a')) \setminus a, l'(a'))\}$$
  
 $\cup \{a \to \ell \in \Gamma' \times 2^{AP} \mid a \text{ contains } a \text{ cycle and } \ell = l'(a)\}$   
 $\cup \{\ell \to \ell \mid \ell \in 2^{AP}\}$ 

4.  $S'_0 = \{a_{s_0} = \text{ReachF}(\{s_0\}, l(s_0)|_{AP}) \mid s_0 \in S_0\}.$ 

The above Definition details how to build a SOG  $\widehat{\mathcal{K}}_{AP}$ . The set  $\mathcal{S}'_0$  of initial states of  $\widehat{\mathcal{K}}_{AP}$  is composed by the set of homogenous aggregates  $a_{s_0}$  satisfying  $a_{s_0} = \operatorname{ReachF}(\{s_0\}, l(s_0)_{|AP})$ , i.e., for each initial state  $s_0$  of  $\mathcal{K}$ ,  $a_{s_0}$  is the set of reachable states s' from  $s_0$  in  $\mathcal{K}$  such that  $l(s') \stackrel{AP}{=} l(s_0)$ . The set  $\mathcal{S}'$  of states of  $\widehat{\mathcal{K}}_{AP}$  is composed of two kinds of nodes:

- 1. homogenous aggregates *a* satisfying  $a = \text{ReachF}(a, \lambda_{AP}(a))$
- 2. divergent states  $\ell \in 2^{AP}$  labeled with subsets of atomic propositions of AP.

For the transition relation of  $\widehat{\mathcal{K}}_{AP}$ ,  $\mathcal{R}'$  is composed of three kinds of edges:

- 1. case *a* and *a'* are two aggregates of  $\Gamma': a \to a' \in \mathcal{R}'$  iff  $l_{AP}(a) \neq l_{AP}(a')$  and *a'* contains every state  $s' \in \Gamma$  satisfying  $l(s')_{|AP} = l'(a')$  and  $s \to s' \in \mathcal{R}$  with  $s \in (a \cup a')$ .
- 2. case *a* is an aggregates of  $\Gamma'$  and  $\ell$  is a divergent state:  $a \to \ell$  iff *a* contains a cycle and  $l'(a) = \ell$ .
- 3. each divergent state  $\ell \in 2^{AP}$  has a self-loop  $\ell \to \ell$ .

Figure 7.1d shows the SOG  $\widehat{\mathcal{K}}_{\{a,b\}}$  built from the Kripke structure  $\mathcal{K}$  of Figure 7.1b according to  $AP = \{a, b\}$  (ignoring the atomic proposition *c* of  $\mathcal{K}$ ).

The initial state of  $\widehat{\mathcal{K}}_{\{a,b\}}$  is an aggregate  $\{s_0, s_1, s_2, s_3\}$  because they agree on the value of atomic propositions observed in  $AP = \{a,b\}$ :  $l(s_0)_{|\{a,b\}} = l(s_1)_{|\{a,b\}} = l(s_2)_{|\{a,b\}} = l(s_3)_{|\{a,b\}} = a\overline{b}$ . This initial aggregate contains a cycle so one of its successors is a divergent state labeled by  $a\overline{b}$ .

The constructed SOG  $\widehat{\mathcal{K}}_{\{a,b\}}$  is also a Kripke structure, that allows to check any stutterinvariant property over the alphabet  $2^{\{a,b\}}$ . As example, Figure 7.1e represents  $\mathcal{G} \otimes \widehat{\mathcal{K}}_{\{a,b\}}$ , the synchronous product between the TGBA  $\mathcal{G}$  of  $a \cup b$  and the SOG  $\widehat{\mathcal{K}}_{\{a,b\}}$ .

**Theorem 2** ([58]). *Given a Kripke Structure*  $\mathcal{K}$  *defined on*  $AP_{\mathcal{K}}$ *, then the SOG*  $\widehat{\mathcal{K}}_{AP}$  *of*  $\mathcal{K}$  *built over*  $AP \subseteq AP_{\mathcal{K}}$  preserves any stuttering-invariant property  $\mathcal{A}$  on AP. In other words:  $\mathscr{L}(\mathcal{A} \otimes \mathcal{K}) \neq \emptyset \iff \mathscr{L}(\mathcal{A} \otimes \widehat{\mathcal{K}}_{AP}) \neq \emptyset$ .

#### 7.3.2 SOG for TGTA (SOG-TGTA)

Klai and Poitrenaud [58] proposed another variant of SOG that does not use divergent states. In a TGTA-based approach, we chose to use this variant of SOG, and we call it SOG-TGTA in this work. Instead of using divergent states, this SOG-TGTA has a self-loop on each aggregate that contains a cycle [58]. In addition, in TGTA all stuttering transitions are self-loops. Therefore, the synchronization between this stuttering self-loops and the self-loops of SOG-TGTA only produces self-loops in the product automaton, and therefore does not generate new states in this product.

**Definition 45** (Symbolic Observation Graph [37] (SOG-TGTA)). Let  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$  be a Kripke structure over  $AP_{\mathcal{K}}$ . A SOG-TGTA over  $AP \subseteq AP_{\mathcal{K}}$  of  $\mathcal{K}$  is defined as  $\widehat{\mathcal{K}}'_{AP} = \langle S', S'_0, \mathcal{R}', l' \rangle$  a Kripke structure over AP satisfying :

- a Kripke structure over  $\overrightarrow{AP}$  satisfying : 1.  $S' = \left\{ a \in 2^{S} \setminus \{\emptyset\} \middle| \begin{array}{c} a \text{ is homogeneous w.r.t. } AP \\ a = \operatorname{ReachF}(a, l_{AP}(a)) \end{array} \right\}$ 
  - 2.  $\forall a \in \mathcal{S}', l'(a) = l_{AP}(a)$
  - 3.  $\mathcal{R}' = \{a \to a' \in \mathcal{S}' \times \mathcal{S}' \mid a' = \text{ReachF}(\text{SuccF}(a, l'(a')) \setminus a, l'(a'))\}$  $\cup \{a \to a \in \Gamma' \times \Gamma' \mid a \text{ contains } a \text{ cycle } \}$
  - 4.  $S'_0 = \{ \operatorname{ReachF}(\{s_0\}, l(s_0)|_{AP}) \mid s_0 \in S_0 \}.$

The only difference between the above Definition 45 and the Definition 44 of SOG is the fact that SOG-TGTA does not use the divergent states labeled with the elements of  $2^{AP}$ . These divergent states are replaced in SOG-TGTA by adding a self-loop on each aggregate that contains a cycle (see point 3 of Definition 45). The obtained SOG-TGTA contains only one kind of nodes: homogenous aggregates.

**Theorem 3** ([58]). Given a Kripke Structure  $\mathcal{K}$  defined on  $AP_{\mathcal{K}}$ , then the SOG-TGTA  $\widehat{\mathcal{K}}'_{AP}$  of  $\mathcal{K}$  built over  $AP \subseteq AP_{\mathcal{K}}$  preserves any stuttering-invariant property  $\mathcal{A}$  on AP. In other words:  $\mathscr{L}(\mathcal{A} \otimes \mathcal{K}) \neq \emptyset \iff \mathscr{L}(\mathcal{A} \otimes \widehat{\mathcal{K}}'_{AP}) \neq \emptyset$ .

Figure 7.2d shows an example of a SOG-TGTA  $\mathcal{K}'_{\{a,b\}}$  and Figure 7.2e presents the product  $\mathcal{T} \otimes \widehat{\mathcal{K}}'_{\{a,b\}}$  of  $\mathcal{K}'_{\{a,b\}}$  with the TGTA  $\mathcal{T}$  of  $a \cup b$ .  $\widehat{\mathcal{K}}'_{\{a,b\}}$  is similar to the SOG  $\widehat{\mathcal{K}}_{\{a,b\}}$  of Figure 7.1d, but without the divergent state labeled with  $a\bar{b}$ . In addition, the initial aggregate of  $\mathcal{K}'_{\{a,b\}}$  has a self-loop because it contains a cycle.

We can notice that  $\mathcal{T} \otimes \widehat{\mathcal{K}}'_{\{a,b\}}$  is smaller than the product  $\mathcal{G} \otimes \widehat{\mathcal{K}}_{\{a,b\}}$  (Figure 7.1e) using the TGBA  $\mathcal{G}$  (recognizing the same formula  $a \cup b$  as the TGTA  $\mathcal{T}$ ). We will present in Section 7.6 an experimental comparison that will confirm this observation.

# 7.4 Symbolic Observation Product (SOP)

The *Symbolic Observation Product* (SOP) [37] is a dynamic extension of SOG that exploits the fact that the number of observed atomic propositions decreases as we progress in the property automaton. The goal of this extension is to allow further aggregation than SOG. However, contrary to SOG, which is an abstraction of a Kripke structure, SOP is an hybrid synchronous product between a Kripke structure and the TGBA of a stutter-invariant property. In this section, we start by giving the definition and an illustrative example of SOP. Then, we present SOP-TGTA, an adaptation of SOP that uses TGTA instead of TGBA to represent the property automaton.

In Section 7.6.4, we will present the results of an experimental comparison between the original SOP approach based on TGBA and our SOP-TGTA approach based on TGTA.

Given a TGBA  $\mathcal{G} = \langle Q, I, \delta, \mathcal{F} \rangle$  or a TGTA  $\mathcal{G} = \langle Q, I, U, \delta, \mathcal{F} \rangle$ , the alphabet FV(q) of a state  $q \in Q$  is defined as the union of the atomic propositions which can be observed from q. Formally,  $FV(q) = \bigcup_{q_1 \stackrel{\phi, F}{\longrightarrow} q_2 \in \delta^*(q)} FV(\phi)$  with  $\delta^*(q)$  is the set of transitions reachable from a state q. From this definition of FV(q), we can easily deduce that for any  $q_1 \stackrel{\phi, F}{\longrightarrow} q_2 \in \delta$ , we have  $FV(q_1) \supseteq FV(q_2)$ . In other words, the set of observed atomic propositions decreases as we progress through the successive states of the property automaton (TGBA or TGTA).

#### 7.4.1 SOP

SOP [38] is an hybrid product that is constructed over a dynamic alphabet (FV) which decreases as the construction of the product progresses. This allows to obtain larger aggregates and therefore fewer number of states than in SOG approach.

**Definition 46** (SOP of a TGBA and a Kripke structure). *Given a stutter-invariant TGBA*  $\mathcal{G} = \langle Q, I, \delta, \mathcal{F} \rangle$  over AP and a Kripke structure  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$ , the Symbolic Observation Product of  $\mathcal{G}$  and  $\mathcal{K}$  is the TGBA denoted  $\mathcal{G} \otimes \mathcal{K} = \langle Q_{\widehat{\otimes}}, I_{\widehat{\otimes}}, \delta_{\widehat{\otimes}}, \mathcal{F} \rangle$  where:

1. 
$$Q_{\widehat{\otimes}} = Q' \cup D'$$
 where states constructed from aggregates are in  $Q'$  and divergent states are  
in  $D'$ :  
 $Q' = \begin{cases} (a, a) \in Q \times (2^{\Gamma} \setminus \{\emptyset\}) \\ a \text{ is homogeneous w.r.t. FV}(q) \end{cases}$ 

$$\begin{aligned} \mathcal{Q}' &= \left\{ \begin{array}{l} (q,a) \in \mathcal{Q} \times (2^{1} \setminus \{\emptyset\}) \middle| a = \operatorname{ReachF}(a, l_{\operatorname{FV}(q)}(a)) \\ \mathcal{D}' &= \{(q,\ell) \mid q \in \mathcal{Q} \text{ and } \ell \in 2^{\operatorname{FV}(q)} \} \end{aligned} \right. \end{aligned}$$

$$2. \quad \delta_{\widehat{\otimes}} = \begin{cases} (q_1, a_1) \xrightarrow{\ell, F} (q_2, a_2) & | (q_1, a_1) \in Q', (q_2, a_2) \in Q', \ell = l_{\mathrm{FV}(q_1)}(a_1), \\ \exists \phi \in \mathbb{B}(AP) \ s.t. \ q_1 \xrightarrow{\phi, F} q_2 \in \delta, \ and \ \ell \models \phi, \\ \exists \ell' \in 2^{\mathrm{FV}(q_2)} \ s.t. \ a_2 = \mathrm{ReachF}(\mathrm{SuccF}(a_1, \ell') \setminus a_1, \ell') \end{cases} \\ \cup \begin{cases} (q_1, a) \xrightarrow{\ell_1, F} (q_2, \ell_2) \\ a \ contains \ a \ cycle, \ \ell_2 = \ell_{1|\mathrm{FV}(q_2)}, \\ \exists \phi \in \mathbb{B}(AP) \ s.t. \ q_1 \xrightarrow{\phi, F} q_2 \in \delta, \ and \ \ell_1 \models \phi \end{cases} \\ \cup \begin{cases} (q_1, \ell_1) \xrightarrow{\ell_1, F} (q_2, \ell_2) \\ \exists \phi \in \mathbb{B}(AP) \ s.t. \ q_1 \xrightarrow{\phi, F} q_2 \in \delta, \ and \ \ell_1 \models \phi \end{cases} \end{cases}$$

3.  $I_{\widehat{\otimes}} = \{(q_0, \operatorname{ReachF}(\{s_0\}, l(s_0)_{|\operatorname{FV}(q_0)})) \mid (q_0, s_0) \in I \times S_0\}$ We have  $\mathscr{L}(\mathcal{G} \otimes \mathcal{K}) \neq \emptyset \iff \mathscr{L}(\mathcal{G} \widehat{\otimes} \mathcal{K}) \neq \emptyset$  by construction.

As in SOG, the set of states  $Q_{\otimes}$  of a SOP is composed of two kinds of states in Q' and D'. The states of Q' are pairs of the form (q, a), where q is a state of TGBA and a is an aggregate of states from the Kripke structure. a is similar to an aggregate in a SOG but computed according to the alphabet FV(q) instead of all AP. The states of D' are the divergent states of the SOP.

The SOP transition relation  $\delta_{\widehat{\otimes}}$  is composed of three parts. The first part contains the transitions of the form  $(q_1, a_1) \rightarrow (q_2, a_2)$ , where  $q_1, q_2$  are two successive states of the TGBA, and  $a_1, a_2$  are two aggregates of states from the Kripke structure. The aggregate  $a_2$  contains the successors of states of  $a_1$  and its computation is similar to SOG, except that  $a_2$  is homogeneous w.r.t. the set  $FV(q_2)$  instead of all *AP*. The second and third parts of the SOP transition relation  $\delta_{\widehat{\otimes}}$  are also similar to the SOG transition relation parts for cycle detection. Figure 7.1f shows an example of a SOP  $\mathcal{G} \otimes \mathcal{K}$  computed from the Kripke structure  $\mathcal{K}$  and the TGBA  $\mathcal{G}$  of  $a \cup b$ . The difference between the SOP and the product  $\mathcal{G} \otimes \widehat{\mathcal{K}}_{\{a,b\}}$  using the SOG (Figure 7.1e), is mainly when the TGBA  $\mathcal{G}$  reaches the state  $q_1$ . Indeed, from this state, the alphabet  $FV(q_1)$  becomes empty. This allows the SOP to aggregate the states  $\{s_4, s_5, s_6, s_7\}$ . In addition, these states form a cycle in  $\mathcal{K}$ , and therefore implies to add a divergent state  $(q_1, \top)$  in the SOP.

#### 7.4.2 SOP Using TGTA (SOP-TGTA)

In this section, we propose a TGTA-based SOP, called SOP-TGTA. The main difference between SOP and SOP-TGTA is related to the changesets labeling the TGTA and their synchronization with the states of the Kripke structure.

**Definition 47** (SOP of a TGTA and a Kripke structure). *Given a TGTA*  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F} \rangle$  over *AP and a Kripke structure*  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$ , the SOP-TGTA of  $\mathcal{T}$  and  $\mathcal{K}$  is the TGTA denoted  $\mathcal{T} \otimes \mathcal{K} = \langle Q_{\widehat{\otimes}}, I_{\widehat{\otimes}}, U_{\widehat{\otimes}}, \delta_{\widehat{\otimes}}, \mathcal{F} \rangle$  where:

1.  $Q_{\widehat{\otimes}} = Q' \cup D'$  where states of the automaton are synchronized with aggregates in Q' and with divergent states in D':

$$\begin{aligned} \mathcal{Q}' &= \left\{ (q,a) \in \mathcal{Q} \times (2^{\mathcal{S}} \setminus \{ \emptyset \}) \middle| \begin{array}{l} a \text{ is homogeneous w.r.t. } \mathrm{FV}(q) \\ a &= \mathrm{ReachF}(a, l_{\mathrm{FV}(q)}(a)) \end{aligned} \right\} \\ \mathcal{D}' &= \{ (q,\ell) \mid q \in \mathcal{Q} \text{ and } \ell \in 2^{\mathrm{FV}(q)} \} \end{aligned}$$

$$2. \quad \delta_{\widehat{\otimes}} = \begin{cases} (q_1, a_1) \xrightarrow{(\ell \oplus \ell'), F} (q_2, a_2) & | (q_1, a_1) \in Q', (q_2, a_2) \in Q', \ell = l_{FV(q_1)}(a_1), \\ \exists \phi \in \mathbb{B}(AP) \ s.t. \ q_1 \xrightarrow{\phi, F} q_2 \in \delta, \\ \exists \ell' \in 2^{FV(q_1)} \ s.t. \ (\ell \oplus \ell') \models \phi, \ and \\ a_2 = \operatorname{ReachF}(\operatorname{SuccF}(a_1, \ell'_{|FV(q_2)}) \setminus a_1, \ell'_{|FV(q_2)}) \end{cases} \\ \cup \begin{cases} (q_1, a) \xrightarrow{\emptyset, F} (q_2, \ell_2) \\ \exists \phi \in \mathbb{B}(AP) \ s.t. \ q_1 \xrightarrow{\phi, F} q_2 \in \delta, \ and \ \emptyset \models \phi \end{cases} \\ \cup \begin{cases} (q_1, \ell_1) \xrightarrow{\emptyset, F} (q_2, \ell_2) \\ \exists \phi \in \mathbb{B}(AP) \ s.t. \ q_1 \xrightarrow{\phi, F} q_2 \in \delta, \ and \ \emptyset \models \phi \end{cases} \end{cases}$$

3. 
$$I_{\widehat{\otimes}} = \{(q_0, \operatorname{ReachF}(\{s_0\}, l(s_0)_{|\operatorname{FV}(q_0)})) \mid (q_0, s_0) \in I \times S_0 \text{ and } l(s_0) \models U(q_0)\},\$$

4.  $\forall (q_0, a_0) \in I_{\widehat{\otimes}}, U_{\widehat{\otimes}}((q_0, a_0)) = l_{\mathrm{FV}(q_0)}(a_0).$ We have  $\mathscr{L}(\mathcal{T} \otimes \mathcal{K}) \neq \emptyset \iff \mathscr{L}(\mathcal{T} \widehat{\otimes} \mathcal{K}) \neq \emptyset$  by construction.

The set of states  $Q_{\otimes}$  of a SOP-TGTA is the same as in SOP. But, the transition relation is a little different because it is based on changesets. As in SOP, the transition relation  $\delta_{\otimes}$  of SOP-TGTA is composed of three rules. The first rule defines the transitions between aggregates, the second rule is about transitions between aggregates and divergent states, and the third rule is for transitions between divergent states.

The transitions between aggregates are of the form  $(q_1, a_1) \rightarrow (q_2, a_2)$ , where  $q_1, q_2$  are two successive states of the TGTA, and  $a_1, a_2$  are two aggregates of states from the Kripke structure. Each aggregate  $a_2$  contains the successors of states of  $a_1$ , filtered to keep only those satisfying a valuation  $\ell'_{|FV(q_2)}$ , where  $\ell'$  satisfies  $(\ell \oplus \ell') \models \phi$  with  $\ell = l_{FV(q_1)}(a_1)$  and  $q_1 \stackrel{\phi, F}{\longrightarrow} q_2$ . In other words,  $\ell'$  is obtained from  $\ell = l_{FV(q_1)}(a_1)$  by applying one changeset from the set of changesets encoded by  $\phi$ . In addition, each aggregate  $a_2$  is homogeneous w.r.t. the set of atomic prositions of FV( $q_2$ ). In the second and third parts of  $\delta_{\hat{\otimes}}$  of SOP-TGTA, all the transitions between aggregates and divergent states and all the self-loops on divergent states are labeled with an empty changeset  $\emptyset$ . Figure 7.2f shows an example of a SOP-TGTA  $\mathcal{T} \hat{\otimes} \mathcal{K}$  computed from the Kripke structure  $\mathcal{K}$ and the TGTA  $\mathcal{T}$  of  $a \cup b$ . The first difference between the SOP-TGTA  $\mathcal{T} \hat{\otimes} \mathcal{K}$  and the product  $\mathcal{T} \otimes \hat{\mathcal{K}}'_{\{a,b\}}$  using SOG-TGTA (Figure 7.2e), is the divergent state labeled with  $(q_0, a\bar{b})$ . This divergent state is added to the SOP-TGTA because the aggregate  $\{s_0, s_1, s_2, s_3\}$  of the initial state contains a cycle. The second difference between  $\mathcal{T} \otimes \mathcal{K}$  and  $\mathcal{T} \otimes \widehat{\mathcal{K}}'_{\{a,b\}}$  comes from the fact that the alphabet  $FV(q_1)$  becomes empty when the TGTA  $\mathcal{T}$  reaches the state  $q_1$ . This allow the SOP-TGTA to aggregate the states  $\{s_4, s_5, s_6, s_7\}$  of  $\mathcal{K}$ . In addition, this aggregate contains a cycle, and therefore allows to add in  $\mathcal{T} \otimes \mathcal{K}$  a divergent state labeled with  $(q_1, \top)$ .

In our example of verification of the LTL property  $a \cup b$  on the Kripke structure  $\mathcal{K}$ , we observe that the SOP-TGTA  $\mathcal{T} \otimes \mathcal{K}$  of Figure 7.2f is smaller than the original SOP  $\mathcal{G} \otimes \mathcal{K}$  of Figure 7.1f.

# 7.5 Self-Loop Aggregation Product (SLAP)

SLAP [38] is an hybrid synchronous product, in which the aggregation of Kripke structure states is based on the self-loops of the property automaton. This section presents the original SLAP based on TGBA and SLAP-TGTA, a variant of SLAP based on TGTA. In Section 7.6.5, we will show the results of an experiment comparing the performance of SLAP and SLAP-TGTA.

**Definition 48.** Given a TGBA  $\mathcal{G} = \langle Q, I, \delta, \mathcal{F} \rangle$  or a TGTA  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F} \rangle$ , for a state  $q \in Q$ , SF(q) encodes the Self-loop Formulas labeling edges  $q \rightarrow q$ . Formally,

$$\mathrm{SF}(q) = \bigvee_{q \xrightarrow{\phi, F} q \in \delta} \phi$$

Similar to the symbolic operations SuccF and ReachF defined in Section 7.3, for SLAP approach, we define two other operations FReach $(a,\phi)$  and Succ $\oplus(a,\phi)$  with  $a \subseteq S$  is a set of states of  $\mathcal{K}$  and  $\phi \in \mathbb{B}(AP)$  is a propositional formula, such that:

- FSucc(a, φ) = {s' ∈ S | ∃s ∈ a, R(s, s') ∧ ∃ℓ, [ℓ ⊨ φ ∧ L(s, ℓ)]}, i.e, Filter a to only keep states satisfying φ, then produce their Successors. The difference between SuccF and FSucc is whether the filter is applied on the source or destination states.
- FReach $(a, \phi)$  computes the least subset of S satisfying:
  - $a \subseteq \operatorname{FReach}(a, \phi)$ ,
  - FSucc(FReach $(a, \phi), \phi) \subseteq$  FReach $(a, \phi)$ .

#### 7.5.1 SLAP

A SLAP is an hybrid product between a Kripke structure  $\mathcal{K}$  and a TGBA  $\mathcal{G}$ . The states of the SLAP are pairs of the form (q,a) composed of a state q of  $\mathcal{G}$  and an aggregate a containing successive states of  $\mathcal{K}$  aggregated as long as they model SF(q). These aggregates are computed as symbolic least fixed-points using the operations FSucc and FReach defined in Section 7.2.

**Definition 49** (SLAP of a TGBA and a Kripke structure). *Given a TGBA*  $\mathcal{G} = \langle Q, I, \delta, \mathcal{F} \rangle$  *over AP and a Kripke structure*  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$ , *the* Self-Loop Aggregation Product *of*  $\mathcal{G}$  *and*  $\mathcal{K}$  *is the TGBA denoted*  $\mathcal{G} \boxtimes \mathcal{K} = \langle Q_{\boxtimes}, I_{\boxtimes}, \delta_{\boxtimes}, \mathcal{F} \rangle$  *where:* 

•  $Q_{\boxtimes} = Q \times (2^{\mathcal{S}} \setminus \{\emptyset\})$ 

• 
$$\delta_{\boxtimes} = \begin{cases} (q_1, a_1) \xrightarrow{\top, F} (q_2, a_2) \\ \end{array} \begin{vmatrix} \exists \phi \in \mathbb{B}(AP) \text{ s.t. } q_1 \xrightarrow{\phi, F} q_2 \in \delta, \\ q_1 = q_2 \Rightarrow F \neq \emptyset, \text{ and} \\ a_2 = \operatorname{FReach}(\operatorname{FSucc}(a_1, \phi), \operatorname{SF}(q_2)) \end{cases}$$

•  $I_{\boxtimes} = \{(q^0, \operatorname{FReach}(\{s_0\}, \operatorname{SF}(q^0))) \mid (q_0, s_0) \in I \times S_0\}$ *We have*  $\mathscr{L}(G \otimes \mathscr{K}) \neq \emptyset \iff \mathscr{L}(G \boxtimes \mathscr{K}) \neq \emptyset$  *by construction.* 

All transitions of the constructed SLAP are labeled with the formula  $\top$ , because these labels are irrelevant when checking language emptiness of SLAP.

Figure 7.1g shows an example of SLAP  $\mathcal{G} \boxtimes \mathcal{K}$  obtained from the Kripke structure  $\mathcal{K}$  and the TGBA  $\mathcal{G}$  of  $a \cup b$ . The initial state of  $\mathcal{G} \boxtimes \mathcal{K}$  is the pair  $\left(q_0, \left\{\begin{smallmatrix}s_0 & s_1\\s_2 & s_3\end{smallmatrix}\right\}\right)$ , where  $q_0$  is the initial state of  $\mathcal{G}$  and the aggregate  $a_1 = \left\{\begin{smallmatrix}s_0 & s_1\\s_2 & s_3\end{smallmatrix}\right\}$  is obtained by iteratively aggregating the successors of the states that satisfy  $SF(q_0) = a\bar{b}$ . Then, from the initial state  $(q_0, a_1)$ , we explore the transition  $q^0 \xrightarrow{b} q_1$  of TGBA and we obtain only one successor  $(q_1, a_2)$  with the aggregate  $a_2 = FReach(FSucc(a_1, b), SF(q_1))$  computed as follows:

- The set FSucc(*a*<sub>1</sub>,*b*) contains only the successors of {*s*<sub>4</sub>} because only the state *s*<sub>4</sub> in *a*<sub>1</sub> satisfies *b*, thus FSucc(*a*<sub>1</sub>,*b*) = {*s*<sub>5</sub>};
- SF $(q_1) = \top$  because  $q_1 \xrightarrow{\top} q_1$
- We deduce that a<sub>2</sub> = FReach({s<sub>5</sub>}, ⊤). a<sub>2</sub> contains all the reachable states from s<sub>5</sub> (satisfying ⊤), thus a<sub>2</sub> = {s<sub>4</sub>, s<sub>5</sub>, s<sub>6</sub>, s<sub>7</sub>}.

Finally, we explore the TGBA transition  $q_1 \xrightarrow{\top, \bullet} q_1$  and we obtain an accepting self-loop on  $(q_1, a_2)$  because the successor aggregate of  $a_2$  is  $a_2$  (i.e, FReach(FSucc $(a_2, \top)$ , SF $(q_1)$ ) =  $a_2$ ).

#### 7.5.2 SLAP Using TGTA (SLAP-TGTA)

The SLAP-TGTA is a variant of SLAP based on TGTA instead of TGBA. In SLAP-TGTA, the states of the Kripke structure are aggregated according to the changesets labelling the TGTA transitions. In particular, each SF(q) represents the set of changesets encoded by the Self-loop Formulas labeling edges  $q \rightarrow q$  of the TGTA. Therefore, in SLAP-TGTA the successive states of the Kripke structure are aggregated as long as they change according to the changesets encoded by SF(q). These aggregates are computed as least fixed-points based on changesets using the symbolic operations Succ<sup> $\oplus$ </sup> and Reach<sup> $\oplus$ </sup> defined as follows:

Let  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$  a Kripke structure, encoded by a changeset-based symbolic Kripke structure  $K^{\oplus} = \langle S_0, R^{\oplus}, L \rangle$  (Definition 37 in page 118). For a set of states  $a \subseteq S$  and a propositional formula  $\phi \in \mathbb{B}(AP)$ , we define the following symbolic operations:

- Succ<sup>⊕</sup>(a, φ) = {s' ∈ S | ∃s ∈ a, ∃k, [k ⊨ φ ∧ R<sup>⊕</sup>(s, k, s')]}, i.e, the set of the Successors states of *a* Filtered to keep only those satisfying k ⊨ φ where k = l(s) ⊕ l(s') is a changeset between l(s) and l(s').
- Reach<sup> $\oplus$ </sup>( $a, \phi$ ) computes the least subset of S satisfying:
  - $a \subseteq \operatorname{Reach}^{\oplus}(a, \phi)$ ,
  - Succ<sup> $\oplus$ </sup>(Reach<sup> $\oplus$ </sup>(*a*, $\phi$ ), $\phi$ )  $\subseteq$  Reach<sup> $\oplus$ </sup>(*a*, $\phi$ ).

**Definition 50** (SLAP of a TGTA and a Kripke structure). *Given a TGTA*  $\mathcal{T} = \langle Q, I, U, \delta, \mathcal{F} \rangle$  *over AP and a Kripke structure*  $\mathcal{K} = \langle S, S_0, \mathcal{R}, l \rangle$ , *the* SLAP-TGTA *of*  $\mathcal{T}$  *and*  $\mathcal{K}$  *is the TGTA denoted*  $\mathcal{T} \boxtimes \mathcal{K} = \langle Q_{\boxtimes}, I_{\boxtimes}, \delta_{\boxtimes}, \mathcal{F} \rangle$  *where:* 

•  $Q_{\bowtie} = Q \times (2^{\mathcal{S}} \setminus \{\emptyset\})$ 

• 
$$\delta_{\boxtimes} = \begin{cases} (q_1, a_1) \xrightarrow{\top, F} (q_2, a_2) \\ q_1 = q_2 \Rightarrow F \neq \emptyset, and \\ a_2 = \operatorname{Reach}^{\oplus}(\operatorname{Succ}^{\oplus}(a_1, \phi), \operatorname{SF}(q_2)) \end{cases}$$

•  $q^0_{\boxtimes} = \{(q^0, \operatorname{Reach}^{\oplus}(\{s_0\}, \operatorname{SF}(q^0))) \mid (q_0, s_0) \in I \times S_0 \text{ and } l(s_0) \models U(q_0)\}$ We have  $\mathscr{L}(\mathcal{T} \otimes \mathcal{K}) \neq \emptyset \iff \mathscr{L}(\mathcal{T} \boxtimes \mathcal{K}) \neq \emptyset$  by construction.

For the same reason as in SLAP, the SLAP-TGTA transitions are only labeled with  $\top$ . The reachable states of a SLAP-TGTA are of the form (q, a) where q is a state of the TGTA and a is an aggregate of states of the Kripke structure such that: For each state  $s \in a$ , if s' is a successor of s in the Kripke structure with  $l(s) \oplus l(s') \models SF(q)$ , then  $s' \in a$ .

Figure 7.2g presents  $\mathcal{T} \boxtimes \mathcal{K}$ , an example of SLAP-TGTA computed from the Kripke structure  $\mathcal{K}$  and the TGTA  $\mathcal{T}$  of  $a \cup b$ . Because  $l(s_0) = a\bar{b}c \models U(q_0) = a\bar{b}$ , the initial state of  $\mathcal{T} \boxtimes \mathcal{K}$  is the pair  $(q_0, a_1)$ , where  $a_1$  is computed from  $s_0$  by iteratively aggregates successors that change according to a changeset belonging to (the set of changesets encoded by)  $SF(q_0)$ . Formally,  $a_1 = \text{Reach}^{\oplus}(\{s_0\}, SF(q^0)) = \{s_0 s_1 s_1\}$  (because  $l(s_0) \oplus l(s_1) = l(s_1) \oplus l(s_2) =$  $l(s_2) \oplus l(s_3) = \emptyset \models SF(q_0) = \bar{a}\bar{b}$ ). Then, in order to compute the successors of  $(q_0, a_1)$ , we explore the transition  $q^0 \xrightarrow{b} q_1$  of TGTA. We obtain only one successor  $(q_1, a_2)$  with the aggregate  $a_2 = \text{Reach}^{\oplus}(\text{Succ}^{\oplus}(a_1, b), SF(q_1))$  is computed as follows:

- Succ<sup>⊕</sup>(a<sub>1</sub>,b) = {s<sub>4</sub>} because s<sub>4</sub> is a successor of s<sub>0</sub> ∈ a<sub>1</sub> and it is the unique successor of states of a<sub>1</sub> that satisfies l(s<sub>0</sub>) ⊕ l(s<sub>4</sub>) ⊨ b (l(s<sub>0</sub>) ⊕ l(s<sub>4</sub>) = {b,c}),
- SF $(q_1) = \top$  because  $q_1 \xrightarrow{\top} q_1$
- Thus, a<sub>2</sub> = Reach<sup>⊕</sup>({s<sub>4</sub>}, ⊤) = {s<sub>4</sub>, s<sub>5</sub>, s<sub>6</sub>, s<sub>7</sub>} because a<sub>2</sub> contains all the reachable states from s<sub>4</sub> through any changesets.

Finally, we compute the successors of  $(q_1, a_2)$  by exploring the TGTA transition  $q_1 \xrightarrow{\top, \bullet} q_1$ . This TGTA accepting self-loop also generates an accepting self-loop on state  $(q_1, a_2)$  of the SLAP-TGTA. Indeed, the unique successor of  $(q_1, a_2)$  is itself because:

Reach<sup>⊕</sup>(Succ<sup>⊕</sup>( $a_2$ ,  $\top$ ),  $\top$ ) =  $a_2$  (the states of  $a_2$  are in a cycle).

# 7.6 Experimental Comparison of Hybrid Approaches using TGBA vs. TGTA

In the previous sections, we presented three hybrid approaches based on TGBA (SOG, SOP and SLAP) and their variants based on TGTA (SOG-TGTA, SOP-TGTA and SLAP-TGTA). This section presents an experimental evaluation conducted to compare each hybrid approach with its variant based on TGTA.

This experimentation is based on BEEM benchmark [67]. It reuses the same benchmark inputs, formulas and models used to evaluate the fully symbolic approaches in Section 6.4.3 page 124.

#### 7.6.1 Implementation

We have implemented SOG-TGTA, SOP-TGTA and SLAP-TGTA in the same tool LTL-ITS<sup>1</sup> that already contains SOG, SOP, SLAP and the fully symbolic approaches (using TGBA/TGTA). LTL-ITS tool is built on top of libraries<sup>2</sup>: SDD/ITS, Spot, and LTSmin. These three libraries were already presented in the previous chapter, Section 6.4.1, page 122.

The DVE variant of LTSmin [12] is used to produce an ETF file representing the transition relation of each BEEM model. These ETF files are loaded by the SDD/ITS [85] library to encode the symbolic transition relations (R(s,s')) of the Kripke structures, used to implement the symbolic operations ReachF and FReach. For SLAP-TGTA, it is a changeset-based symbolic transition relations  $R^{\oplus}(s,s')$  (Definition 37) that are built from the ETF files. It is used to implement the symbolic operation Reach<sup> $\oplus$ </sup> for the SLAP-TGTA aggregates computation.

The Spot library [36] is used to translate the LTL properties into TGBA or TGTA. It is also used to perform the emptiness check of explicit graphs, such as the synchronous products TGBA/SOG and TGTA/SOG-TGTA. In addition, the hybrid synchronous products SOP, SLAP, SOP-TGTA and SLAP-TGTA are also handled by the emptiness check of Spot. Indeed, they are explicit graphs in which each node stores a set of states encoded as a Decision Diagram. These sets of states are computed using least fixed-points (ReachF, FReach or Reach<sup> $\oplus$ </sup>).

A SOG-TGTA is implemented in the same way as a SOG, as a concrete class of the Kripke structure interface provided by Spot. During the the emptiness check of the products TGBA/SOG and TGTA/SOG-TGTA, the SOG and the SOG-TGTA nodes are constructed on-the-fly using the implementation of the symbolic operation ReachF.

Similar to SOP and SLAP, we have implemented SOP-TGTA and SLAP-TGTA as concrete classes of the synchronous product interface of Spot. During the the emptiness check, the nodes of these four hybrid products are built on-the-fly from the states of the property automaton (TGBA or TGTA) and using the symbolic operations: ReachF for SOP and SOP-TGTA, FReach for SLAP and Reach<sup> $\oplus$ </sup> for SLAP-TGTA.

#### 7.6.2 Results

The results of our experimentations are presented as scatter plots using logarithmic scale. Each scatter plot compares an hybrid approach against its variant based on TGTA. Each point represents the comparison of the performance of the model checking for a model and formula pair. Any process that exceeded 60 minutes of runtime or 6GB of RAM was aborted (thus, the answer of the model checker was not reported for some cases). In our scatter plots, these aborted experiments are plotted as being three times higher than the maximum of the other values. Thus, these points appear separately (by the wide white band) from the other experiments that succeed.

http://ddd.lip6.fr

<sup>&</sup>lt;sup>2</sup>Respectively http://ddd.lip6.fr, http://spot.lip6.fr, and http://fmt.cs.utwente.nl/tools/ ltsmin.





Figure 7.3: Performance comparison of SOG against SOG-TGTA. Left: time (in seconds); Right: memory (in kilobytes).

		SOG	SOG-TGTA
Verified Formulas	Win	315	826
	Fail	39	4
Violated Formulas	Win	1441	609
	Fail	14	5

Table 7.1: On all successful experiments, we count the number of cases a specific method has (Win) the best time. The Fail line shows the cases were an approach failed to solve an experiment solved by the other approach.

The scatter plots of Figure 7.3 compare the performance of two hybrid approaches: the first is based on TGBA and SOG (called just SOG approach in the following); the second is based on TGTA and SOG-TGTA (called SOG-TGTA approach). Each point of the left and right scatter plots compares respectively the time and memory used in the model checking of each pair (formula, model). The x-axis represents the performance of SOG approach and the y-axis shows the performance of SOG-TGTA approach, so the points below the diagonal correspond to cases where the SOG-TGTA approach is better. Symmetrically, the points above the diagonal corresponds to points were the SOG approach is better. The points represented by green squares correspond to verified formulas (empty products), and the black crosses correspond to violated formulas (non-empty products).

In the two scatter plots, we observe that for verified formulas (green points), the SOG-TGTA approach outperforms the SOG approach (in time and memory). This result is similar to the comparison between the explicit approaches based on TGBA versus TGTA, presented in Chapter 5 (Section 5.5.2 page 99). This similarity is justified by the fact that these approaches are based on

traditional explicit synchronous products.

For violated formulas, the SOG approach outperforms the SOG-TGTA approach in the cases where the execution time is less than one second. On the contrary, for hard cases, there are more cases that failed using SOG than using SOG-TGTA (compare the "aligned" black points at top and on right of scatter plots, or for more details see Table 7.1).

In total in these scatter plots, SOG failed in 53 cases solved by SOG-TGTA, while SOG-TGTA approach failed for only 9 cases solved by SOG. In addition, 875 cases are not solved by the two approaches within the time and memory limits. In the experiments that have not failed, SOG-TGTA was at least a ten times faster than SOG in 20 cases, and twice times faster in 287 cases. A contrario, SOG was at least a ten times faster than SOG-TGTA in 4 cases, and twice times faster in 144 cases.

Table 7.1 presents the best and the failed approach for each experiment where at least one approach succeeded. This table does not take into account cases where the two approaches failed.



#### 7.6.4 SOP versus SOP-TGTA

Figure 7.4: Performance comparison of SOP against SOP-TGTA. Left: time (in seconds); Right: memory (in kilobytes).

		SOP	SOP-TGTA
Verified Formulas	Win	915 (61%)	593 (39%)
	Fail	44 (2%)	327 (21%)
Violated Formulas	Win	957 (47%)	1061 (52%)
	Fail	18 (0%)	8 (0%)

Table 7.2: On all successful experiments, we count the number of cases a specific method has (Win) the best time. The Fail line shows the cases were each approach failed to compute the result of an experiment solved by the other approach.



Figure 7.5: Performance comparison of SLAP against SLAP-TGTA. Left: time (in seconds); Right: memory (in kilobytes).

The scatter plots of Figure 7.4 compares the performance of SOP against SOP-TGTA. Each point compare the time (left scatter plot) and memory (right scatter plot) used to perform the model checking of each pair (formula, model) of our benchmark. The points below the diagonal correspond to cases where the SOP-TGTA approach is better, and for the other points the SOP approach is better.

On the one hand, the results of the scatter plots are very difficult to interpret, because there are many points where SOP is better and many others with SOP-TGTA which is the best (SOP-TGTA was at least a hundred times faster than SOP in 36 cases, ten times faster in 160 cases, and twice times faster in 579 cases. A contrario, SOP was at least one hundred times faster than SOP-TGTA in 38 cases, ten times faster in 169 cases, and twice times faster in 511 cases).

On the other hand, if we look at the scatter plots in more detail, we can observe that there are more cases of failure for SOP-TGTA (represented by the linear cloud at the top of the scatter plots) than the SOP approach (linear cloud on the right of the scatter plots). Indeed, SOP-TGTA failed for 335 experiments where SOP reached the result, while SOP only failed for 62 cases solved by SOP-TGTA. Table 7.2 gives more details for these failed cases by distinguishing verified and violated formulas (without taking into account the cases where the two approaches failed). This table also shows the number of cases in which each approach is better.

#### 7.6.5 SLAP versus SLAP-TGTA

The scatter plots of Figure 7.5 compares the performance of SLAP against SLAP-TGTA. The left scatter plot compares the time performance and the right concerns the memory consumption. The points below the diagonal correspond to the cases where SLAP-TGTA is better.

The interpretation of the scatter plots results depends on the colors of the points. For black crosses that correpond to violated formulas, SLAP-TGTA is more efficient than SLAP in most cases. For the green squares representing the verified formulas, the scatter plots are difficult to

		SLAP	SLAP-TGTA
Verified Formulas	Win	981 (50%)	964 (50%)
	Fail	57 (2%)	261 (13%)
Violated Formulas	Win	644 (31%)	1374 (68%)
	Fail	28 (1%)	16 (0%)

Table 7.3: On all successful experiments, we count the number of cases a specific method has (Win) the best time. The Fail line shows the cases were each approach failed to compute the result of an experiment solved by the other approach.

interpret, there are many cases on both sides of the diagonal. In addition, according to Table 7.3, there are roughly the same number of cases where each approach is better than the other. However, Table 7.3 shows that SLAP is clearly better for failed cases. In total, SLAP-TGTA failed in 277 experiments solved by SLAP, while SLAP failed for only 85 cases solved by SLAP-TGTA. In the other cases where the two approaches were successful, we observe a relative advantage for SLAP-TGTA. Indeed, on the one hand, SLAP-TGTA was at least a hundred times faster than SLAP in 85 cases, ten times faster in 646 cases, and twice times faster in 1605 cases. On the other hand, SLAP was at least one hundred times faster than SLAP-TGTA in 12 cases, ten times faster in 95 cases, and twice times faster in 434 cases only.

We believe that the two approaches SLAP and SLAP-TGTA are complementary and very different because of the fact that SLAP-TGTA aggregates are based on changesets and therefore are very different from SLAP aggregates that are based on valuations. Thus, these two approaches can be considered complementary and can be launched in parallel in order to retrieve the result of the fastest approach.

# 7.7 Conclusion

In this chapter, we proposed three hybrid approaches variants using TGTA: SOG-TGTA, SOP-TGTA and SLAP-TGTA.

SOG-TGTA is a variant of SOG without divergent states. They are replaced in SOG-TGTA by adding a self-loop on each aggregate that contains a cycle. Adding these self-loop is better than adding divergent states for SOG-TGTA because in TGTA all stuttering transitions are self-loops, and therefore adding self-loops in SOG-TGTA does not generate new states in the product between TGTA and SOG-TGTA.

SOP-TGTA is an adaptation of SOP to TGTA. The difference between SOP and SOP-TGTA appears when computing the transitions of the hybrid product. Indeed, the synchronisation of the transitions between the Kripke structure aggregates and the TGTA transitions is based on changesets.

SLAP-TGTA is also a adaptation of SLAP to use TGTA instead of TGBA. The two variants (SLPA and SLAP-TGTA) are based on the aggregation of the states of a Kripke structure according to the self-loops of the formula automaton (i.e., TGBA for SLAP and TGTA for SLAP-TGTA). However, the obtained aggregates are very different between the two variants. In a SLAP-TGTA,

the states of an aggregate change according to the changesets of the TGTA self-loops (instead of satisfying the valuations labeling the TGBA self-loops in the case of SLAP).

We implemented and experimentally compared the performance of each TGTA based hybrid approach (SOG-TGTA, SOP-TGTA and SLAP-TGTA) against its reference variant (SOG, SOP and SLAP). The obtained results show that SOG-TGTA is better than SOG to check verified properties but it is worse for the violated properties. For SOP versus SOP-TGTA, the results are harder to interpret. The results seem balanced with a slight advantage for SOP, especially because we have more cases solved by SOP and failed for SOP-TGTA than otherwise. Finally, for SLAP versus SLAP-TGTA, although we can observe an advantage for SLAP-TGTA for violated formulas, the results remain difficult to interpret. There are many cases where SLAP-TGTA is better and vice versa in many cases SLAP is better. In addition, for verified properties, there are many cases that failed for SLAP-TGTA but successful for SLAP. This large difference in performance between SLAP and SLAP-TGTA can be explained by the fact that the aggregates computed using TGBA and TGTA are very different (in SLAP-TGTA, the aggregates are based on changesets instead of valuations as in SLAP).

More generally, for the three hybrid approachs, we believe that the two variants (using TGBA versus TGTA) are complementary, and the best solution is to run the two variants in parallel, then take the result of the faster one.

Our work presented in this chapter about hybrid approaches using TGTA is not finished. In particular, we must look for an optimization that exploits the fact that TGTA is specific to stutter-invariant properties, as we did in the previous chapter about symbolic approaches, in which we have proposed an optimization based on the "stuttering-normalization constraint" of TGTA.

# **Conclusion and Perspectives**

# 8.1 Context

The automata-theoretic approach [89, 90] is traditionally used for the model checking of LTL properties. In this approach, a Kripke structure  $\mathcal{K}_M$  is used to represent the state-space of the model M. The property to check is expressed as an LTL formula  $\varphi$ , then its negation is converted into an  $\omega$ -automaton  $A_{\neg\varphi}$ . The third operation is the synchronization between  $\mathcal{K}_M$  and  $A_{\neg\varphi}$ . This constructs a product automaton  $\mathcal{K}_M \otimes A_{\neg\varphi}$  whose language,  $\mathscr{L}(\mathcal{K}_M) \cap \mathscr{L}(A_{\neg\varphi})$ , is the set of executions of M invalidating  $\varphi$ . The last operation is the emptiness check algorithm that explores the product to tell whether it accepts or not an infinite word, i.e., a counterexample. The model M verifies  $\varphi$  iff  $\mathscr{L}(A_M \otimes A_{\neg\varphi}) = \emptyset$ .

**Problem.** The performance of the emptiness check suffers from the well known state-space explosion problem [87]. The product automaton is often too large to be emptiness checked in a reasonable run time and memory (the size of the product can reach  $(|A_{\neg \varphi}| \times |\mathcal{K}_{\mathcal{M}}|)$  with  $|\mathcal{K}_{\mathcal{M}}|$  often very large).

# 8.2 Existing Work

In order to improve the performance of the LTL model checking, many works have attempted to build optimized property automaton  $A_{\neg\varphi}$ , either by improving the LTL translation, or by proposing several reductions for the automaton produced by this translation such as bisimulation/simulation reductions. In most of this works,  $A_{\neg\varphi}$  is represented using the traditional variant of Büchi Automata with state-based accepting (BA). However, according to [27, 49], moving the accepting conditions from the states to the transitions reduces the size of  $A_{\neg\varphi}$ . For instance, *Transition-based Generalized Büchi Automata* (TGBA) is another variant of Büchi Automata that represent the LTL properties more concisely [49, 36] than BA, because TGBA use *generalized* (i.e., multiple) Büchi acceptance conditions *on transitions* rather than on states.

**Testing Automata (TA).** Hansen et al. [52] propose an alternative to Büchi automata, called Testing Automata (TA), that only represent stutter-invariant properties. According to Geldenhuys and Hansen [46], thanks to their high degree of determinism, the TA allow during the model checking to obtain a smaller product than BA. However, TA have two different ways to accept infinite words (livelock or Büchi), an unfortunate consequence is that the emptiness-check algorithm required must perform two passes on the whole product in the worst case.

**Evaluation of the Testing Automata Approach.** We experimentally evaluated the performance of the model checking approach using TA against two variants of Büchi automata BA and TGBA. In the benchmark results, we distinguished violated formulas (i.e., when a counterexample is found) from verified formulas (i.e., exhibit no counterexample).

For verified formulas, we found that the product reduction achieved by the TA approach was not enough to compensate for the two-pass emptiness check this approach requires. It is therefore better to use the TGBA approach, which is more efficient than TA and BA to prove that a stutterinvariant formula is verified.

For violated formulas, the TA approach usually processes less transitions in the product than the BA approach to find a counterexample. This is especially true on random formulas. With weak-fairness formulas, TGBA are advantaged by their generalized acceptance conditions and are able to beat the TA on the average in half of examples.

## 8.3 Contributions

The general objective of this work is to fight against the state-space explosion problem, by reducing the size of the product automaton and/or by decreasing the amount of time and memory used in the emptiness check of this product.

To achieve this goal, we focus on improving the performance of the model checking for stutterinvariant LTL properties. We firstly extend the work of Geldenhuys and Hansen [46] about Testing Automata (TA), by proposing new types of  $\omega$ -autamata optimized for stutter-invariant LTL properties. These new automata represent all the stuttering-transitions using only self-loops, and (unlike TA) only require a single-pass emptiness check algorithm. Secondly, using our main new type of automata, called *Transition-based Generalized Testing Automata (TGTA)*, we propose contributions to improve the performance of three different approaches, i.e., explicit, symbolic and hybrid model checking approaches, where hybrid means combining explicit and symbolic approaches.

The contributions of this thesis could be summarized as follows:

**Single-pass Testing Automata (STA).** After improving the emptiness check of TA approach to avoid the second pass in particular cases, we propose STA (*Single-pass Testing Automata*), a transformation of TA in a normal form that never requires such a second pass.

An experimental evaluation shows that these improvements compete well on our benchmarks. Especially, STA remain good for violated properties, and also beat TA and TGBA (and BA) in most cases when properties are verified. Unfortunately, we observe in certain cases that the STA increases the size of the product automaton, because the transformation from TA into STA adds an artificial livelock-accepting state in order to remove the second pass.

**Transition-based Generalized Testing Automata (TGTA): A Single-pass and Generalized New Automata.** We propose another type of ω-automata for stutter-invariant properties, better than STA, called TGTA (*Transition-based Generalized Testing Automata*). TGTA combine advantages observed on both TA and TGBA, without the second pass of TA, and without adding an artificial state as in STA.

#### 8.3. Contributions

TGTA inherits from TA the elimination of useless stuttering-transitions, but without introducing a second mode of acceptance (i.e, livelock-accepting states). From TGBA, it inherits the use of generalized acceptance conditions on transitions.

We have run benchmarks to evaluate the performance of the TGTA approach against the TA and TGBA (BA) approaches. The experiments report that, in most cases, TGTA lead to smaller products on the average, and the TGTA approach outperforms the TA and TGBA approaches when the property is verified. However, when the property is violated, the results are difficult to interpret because in this case, the on-the-fly algorithm can stop as soon as it finds a counterexample before exploring the entire product. Changing the order in which non-deterministic transitions of the property automaton are iterated is enough to change the number of states and transitions to be explored before a counterexample is found.

Beyond this ordering luckiness in the case of violated formula, we believe that TGTA is better than TA firstly because TGTA does not require a second pass during the emptiness check . Secondly, in our experiments, we observed that TGTA represent more concisely the LTL formulas using (multiple) generalized acceptance conditions, especially for weak-fairness formulas, for which the number of acceptance conditions is greater than other random formulas.

We also believe that TGTA is better than TGBA because the TGTA construction exploits the fact that it is specific to stutter-invariant properties to remove the useless stuttering-transitions, while the TGBA does not exploit at all this specificity. The constructed TGTA represents all the stuttering-transitions using only self-loops, which can reduce the multiplication of stuttering steps in the product. This advantage of TGTA is better exploited in the following symbolic approach.

**Using TGTA to improve Symbolic Model Checking.** We also use TGTA to improve the symbolic approach.

After showing how to encode a symbolic TGTA, we introduce a new symbolic Kripke structure labeled with changesets on transitions as in TGTA. This changeset-based symbolic Kripke structure simplifies the symbolic transition relation of the product.

The main improvement proposed in this approach is based on the combination of TGTA with saturation technique proposed by Ciardo et al. [20]. We show that the performance of the saturation algorithm greatly benefits from the property of TGTA that all stuttering transitions are self-loops and every state has a stuttering self-loop. In other words, the exploration of stuttering transitions in the product is equivalent to only explore stuttering transitions in the model (remaining in the same TGTA states). This property allowed us to improve the stuttering part in the transition relation of the product, this part is only dependent on the model, and can be evaluated without consulting the transition relation of the TGTA. This allows the saturation algorithm to ignore the symbolic variables encoding the TGTA in this stuttering part of product, and therefore efficiently computes (i.e. saturates) the product nodes corresponding to the variables encoding the model.

We experimentally compared this TGTA-based approach to a symbolic approach based on TGBA and saturation. On our benchmark, using TGTA, we were able to gain one order of magnitude over the TGBA-based approach, for both verified and violated properties.

**Using TGTA in three Hybrid Model Checking Approaches.** In addition of the explicit and symbolic approaches, we evaluate the use of TGTA in the context of hybrid [79, 38] model check-

ing that combines the use of both explicit and symbolic techniques.

We focus on three hybrid techniques proposed in [37]: the *Symbolic Observation Graph* (SOG), the *Symbolic Observation Product* (SOP) and the *Self-Loop Aggregation Product* (SLAP), and for each of them we propose a variant using TGTA: SOG-TGTA, SOP-TGTA and SLAP-TGTA.

Implementation and experimental evaluation of each TGTA-based variant against the corresponding original approach show that the obtained results depend on whether the formula is verified or violated:

- SOG vs. SOG-TGTA: For verified properties, SOG-TGTA outperforms SOG. For violated properties, the SOG approach outperforms the SOG-TGTA approach in the cases where the execution time is low (less than one second). On the contrary, for hard cases, there are more cases that failed using SOG than using SOG-TGTA.
- SOP vs. SOP-TGTA (the results are harder to interpret): SOP and SOP-TGTA are comparable with a slight advantage for SOP, because we have more cases solved by SOP and failed for SOP-TGTA than otherwise.
- SLAP vs. SLAP-TGTA: For violated properties, SLAP-TGTA outperforms SLAP. For verified properties, the results are more difficult to interpret. There are many cases where SLAP-TGTA is better and vice versa in many cases SLAP is better. We believe that SLAP and SLAP-TGTA are complementary, and the best solution is to run the two variants in parallel, then take the result of the faster one.

### 8.4 Perspectives

#### 8.4.1 Improving TGTA-based approaches

Several optimizations can be added to TGTA and TA, such as the simulation-reduction [3], which it is currently only implemented for TGBA in Spot. In addition, the optimizations presented in Section 3.4.3 (page 46) for TA can be easily adapted to TGTA.

Another important optimization is to build on-the-fly the TA and the TGTA during the construction of the synchronous product. Especially when the number of atomic propositions (*AP*) is very large, because this may lead to build a TA or a TGTA with a large number of unnecessary initial states, that are not synchronized with the initial state(s) of the Kripke structure, see for example the product  $\mathcal{K} \otimes \mathcal{T}$  presented in Section 5.6 (page 97). In this example, a TGTA  $\mathcal{T}$  contains 3 initial states and only one of them is synchronized with the initial state the Kripke structure  $\mathcal{K}$ . In addition, the optimizations of Section 3.4.3 (already mentioned above) can be easily integrated in a on-the-fly construction of a TA or a TGTA.

As a future work, an idea would be to provide a direct conversion of LTLX to TGTA, without the intermediate TGBA step. We believe a tableau construction such as the one of Couvreur [27] could be easily adapted to produce TGTA. We can also translate an LTLX formula into an itermediate  $\emptyset$ -TGTA then we apply the stuttering reduction of Property 9 (page 93) to transform the obtained  $\emptyset$ -TGTA into TGTA.

Another idea is to investigate the use of  $\emptyset$ -TGTA to improve the model checking of LTL properties (stutter-invariant or not). Indeed, unlike TGTA, the  $\emptyset$ -TGTA can represent any LTL formula. Furthermore, the procedure proposed in Chapter 5 to build an  $\emptyset$ -TGTA from a TGBA does not exploit the restriction to stutter-invariant properties and therefore can be used to build an  $\emptyset$ -TGTA for any LTL formula.

Finally, our work presented in the last chapter about hybrid approaches using TGTA is not finished. In particular, we must look for an optimization that exploits the fact that TGTA is specific to stutter-invariant properties, as we did in the chapter about the symbolic approach, in which we have proposed an optimization based on the stuttering self-loops of TGTA.

#### 8.4.2 Finding sub-classes of LTL formulas for which TGTA is always efficient

An interesting study for TGTA would be to look for a subclass of LTL formulas for which the product of their TGTA with any model will always be smaller than the product with a TGBA or a BA.

We have already started to look for this type of formulas by analyzing the results of our experiments presented in this work. We believe that the subclass of formulas of the form  $\varphi = FGp$  is a good candidate. In the following, for  $\varphi = FGp$  we illustrate the advantage of TGTA compared to TGBA, in the case of a product with the stuttering parts of a Kripke structure.

Figure 8.1 shows the TGBA and the TGTA for the LTL property  $\varphi = FGp$ .



(a) A TGBA  $\mathcal{A}$  for FG p

(b) A TGTA  $\mathcal{T}$  for FG p

Figure 8.1: A TGBA (left) and a TGTA (right) for the LTL property  $\varphi = FG p$ .



(a) Kripke Structure  $\mathcal{K}_1$  that stutters with p true

(b) Kripke Structure  $\mathcal{K}_2$  that stutters with p false

Figure 8.2: The two possible cases of stuttering parts of a Kripke structure.

In order to compare the products, we analyze the two possible cases of stuttering parts of a Kripke structure, represented in Figure 8.2. The first case of Figure 8.4a is a subpart of Kripke structure that stutters with p true. The products of this subpart with the TGBA  $\mathcal{A}$  and the TGTA

 $\mathcal{T}$  are respectively represented in Figure 8.3 and Figure 8.4. In this case, we can observe that the product using TGTA is smaller than the one using TGBA: the (diagonal) transitions from  $(0, s_n)$ to  $(1, s_{n+1})$  and from  $(0, s_{n+2})$  to  $(0, s_{n+2})$  in the product using TGBA does not exist in the product using TGTA.

In the second case where the Kripke structure stutters with p false (see Figure 8.4b), the products using TGBA and TGTA are the same.

So overall, for the subclass of formulas of the form  $\phi = FG p$ , the stuttering parts of a Kripke structure lead to a smaller parts of product using TGTA than TGBA.



(a) Product of TGBA  $\mathcal{A}$  and Kripke Structure  $\mathcal{K}_1$  that stutters with p true



(b) Product of TGBA  $\mathcal{A}$  and Kripke Structure  $\mathcal{K}_2$  that stutters with p false

Figure 8.3: The two products of the TGBA  $\mathcal{A}$  with the two possible cases of stuttering parts of a Kripke structure.



(a) Product of TGTA  $\mathcal T$  and Kripke Structure  $\mathcal K_{\mathrm{I}}$  that stutters with p true



(b) Product of TGTA  $\mathcal{T}$  and Kripke Structure  $\mathcal{K}_2$  that stutters with p false

Figure 8.4: The two products of TGTA  $\mathcal{T}$  with the two possible cases of stuttering parts of a Kripke structure.

#### 8.4.3 Combining TGTA with Partial Order Reductions

Finally, another future work is to combine the TGTA with other techniques that propose statespace optimizations specific to stutter-invariant properties, such as the partial order reduction implemented in Spin tool [55].

Several partial order reduction techniques have been proposed, as the *stubborn sets* of Valmari [86], the *persistent sets* of Godefroid [50] and the *ample sets* of Peled [69]. The basic idea of these reductions is to prune the state-space  $\mathcal{K}_M$  by identifying equivalent interleaving sequences that only differ by the order of concurrent transitions.

TGTA and partial order reduction are complementary. Indeed, while the TGTA-based approaches focus on optimizing the property automata, the partial order techniques try to reduce the state-space of the model.

# **Bibliography**

- [1] André Arnold. *Finite transition systems. Semantics of communicating systems.* Prentice-Hall, 1994. (Cited on page 13.)
- [2] Tomáš Babiak, Mojmír Křetínský, Vojtěch Řehák, and Jan Strejček. LTL to Büchi automata translation: Fast and more deterministic. In *Proceedings of the 18th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'12)*, volume 7214 of *Lecture Notes in Computer Science*, pages 95–109. Springer, 2012. (Cited on page 30.)
- [3] Tomáš Babiak, Thomas Badie, Alexandre Duret-Lutz, Mojmír Křetínský, and Jan Strejček. 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), volume 7976 of *Lecture Notes in Computer Science*, pages 81–98. Springer, July 2013. (Cited on pages 28, 30, 40, 108, 158 and 173.)
- [4] Christel Baier and Joost-Pieter Katoen. *Principles of model checking*. MIT Press, 2008. ISBN 978-0-262-02649-9. (Cited on page 21.)
- [5] Ala Eddine Ben Salem, Alexandre Duret-Lutz, and Fabrice Kordon. Generalized Büchi automata versus testing automata for model checking. In *Proceedings of the 2nd workshop on Scalable and Usable Model Checking for Petri Nets and other models of Concurrency (SUMo'11)*, volume 726, pages 65–79, Newcastle, UK, June 2011. CEUR. (Cited on page 40.)
- [6] Ala Eddine Ben Salem, Alexandre Duret-Lutz, and Fabrice Kordon. Model Checking using Generalized Testing Automata. *Transactions on Petri Nets and Other Models of Concurrency* (*ToPNoC VI*), 7400:94–122, 2012. (Cited on pages 47, 88 and 122.)
- [7] Ala Eddine Ben Salem, Alexandre Duret-Lutz, Fabrice Kordon, and Yann Thierry-Mieg. Symbolic Model Checking of stutter invariant properties Using Generalized Testing Automata. In 20th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS), volume 8413 of Lecture Notes in Computer Science, pages 440–454, Grenoble, France, April 2014. Springer. (Cited on page 114.)
- [8] B. Bérard, M. Bidoit, A. Finkel, F. Laroussinie, A. Petit, L. Petrucci, and Ph. Schnoebelen. Systems and Software Verification. Model-Checking Techniques and Tools. Springer, 2001. URL http://www.springer.com/computer/swe/book/978-3-540-41523-7. (Cited on page 1.)
- [9] Béatrice Berard, Laure Millet, Maria Potop-Butucaru, Yann Thierry-Mieg, and Sébastien Tixeuil. Formal verification of Mobile Robot Protocols. Technical report, Laboratoire d'Informatique de Paris 6 - LIP6, Laboratory of Information, Network and Communication Sciences - LINCS, Institut Universitaire de France - IUF, May 2013. URL http: //hal.archives-ouvertes.fr/hal-00834061. (Cited on page 13.)

- [10] František Blahoudek, Alexandre Duret-Lutz, Mojmír Křetínský, and Jan Strejček. Is there a best Büchi automaton for explicit model checking? In *Proceedings of the 21th International SPIN Symposium on Model Checking of Software (SPIN'14)*, pages ?-? ACM, July 2014. To appear. (Cited on pages 30 and 40.)
- [11] Roderick Bloem. *Search Techniques and Automata for Symbolic Model Checking*. PhD thesis, University of Colorado, 2001. (Cited on pages 28 and 37.)
- [12] S. C. C. Blom, J. C. van de Pol, and M. Weber. LTSmin: Distributed and Symbolic Reachability. In T. Touili, B. Cook, and P. Jackson, editors, *Computer Aided Verification, Edinburgh*, volume 6174 of *Lecture Notes in Computer Science*, pages 354–359, Berlin, July 2010. Springer Verlag. (Cited on pages 122, 124 and 149.)
- [13] Stefan Blom and Simona Orzan. Distributed branching bisimulation reduction of state spaces. *Electronic Notes in Theoretical Computer Science*, 89(1):99–113, 2003. ISSN 1571-0661. doi: http://dx.doi.org/10.1016/S1571-0661(05)80099-4. URL http://www. sciencedirect.com/science/article/pii/S1571066105800994. {PDMC} 2003, Parallel and Distributed Model Checking (Satellite Workshop of {CAV} '03). (Cited on page 26.)
- [14] Ahmed Bouajjani, Jean-Claude Fernandez, Nicolas Halbwachs, Pascal Raymond, and Christophe Ratel. Minimal state graph generation. *Science of Computer Programming*, 18:247–269, 1992. (Cited on pages 26 and 27.)
- [15] Daniel Brand and Pitro Zafiropulo. On communicating finite-state machines. *Journal of the ACM*, 30(2):323–342, 1983. (Cited on page 13.)
- [16] J. Richard Büchi. On a decision method in restricted second order arithmetic. In *Proceedings* of the International Congress on Logic, Methodology, and Philosophy of Science, Berkley, 1960, pages 1–11. Standford University Press, 1962. (Cited on pages 4 and 19.)
- [17] Jerry R. Burch, Edmund M. Clarke, Kenneth L. McMillan, David L. Dill, and L.J. Hwang. Symbolic model checking: 10<sup>20</sup> states and beyond. In *Proceedings of the Fifth Annual IEEE Symposium on Logic in Computer Science*, pages 1–33, Washington, D.C., 1990. IEEE Computer Society Press. (Cited on pages 3, 7, 12, 114 and 115.)
- [18] Gianfranco Ciardo and Andy Jinqing Yu. Saturation-based symbolic reachability analysis using conjunctive and disjunctive partitioning. In *Proceedings of the 13 IFIP WG 10.5 international conference on Correct Hardware Design and Verification Methods*, volume 3725 of *Lecture Notes in Computer Science*, pages 146–161. Springer-Verlag, 2005. (Cited on pages 119 and 120.)
- [19] Gianfranco Ciardo, Gerarld Lüttgen, and Radu Siminiceanu. Efficient symbolic state-space construction for asynchronous systems. In Mogens Nielsen and Dan Simpson, editors, *Proceedings of the 21st International Conference on Application and Theory of Petri Nets* (ICATPN'00), volume 1825 of Lecture Notes in Computer Science, pages 103–122. Springer-Verlag, 2000. (Cited on page 54.)

- [20] Gianfranco Ciardo, Robert M. Marmorstein, and Radu Siminiceanu. Saturation unbound. In Proceedings of the 9th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'03), volume 2619 of Lecture Notes in Computer Science, pages 379–393. Springer-Verlag, 2003. (Cited on pages 7, 114, 119 and 157.)
- [21] Gianfranco Ciardo, Gerald Lüttgen, and Andrew S. Miner. Exploiting interleaving semantics in symbolic state-space generation. *Formal Methods in System Design*, 31(1):63–100, 2007. (Cited on page 54.)
- [22] Jacek Cichoń, Adam Czubak, and Andrzej Jasiński. Minimal Büchi automata for certain classes of LTL formulas. In *Proceedings of the Fourth International Conference on Dependability of Computer Systems (DEPCOS'09)*, pages 17–24. IEEE Computer Society, 2009. (Cited on page 24.)
- [23] E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. ACM Trans. Program. Lang. Syst., 8(2):244–263, April 1986. ISSN 0164-0925. doi: 10.1145/5397.5399. URL http://doi.acm.org/10.1145/5397.5399. (Cited on page 16.)
- [24] Edmund M. Clarke and Jeannette M. Wing. Formal methods: State of the art and future. ACM Computing Surveys, 28(4):626–643, 1996. (Cited on page 1.)
- [25] Edmund M. Clarke, Orna Grumberg, and Doron A. Peled. *Model Checking*. The MIT Press, 2000. (Cited on page 1.)
- [26] Costas Courcoubetis, Moshe Y. Vardi, Pierre Wolper, and Mihalis Yannakakis. Memoryefficient algorithm for the verification of temporal properties. In Edmund M. Clarke and Robert P. Kurshan, editors, *Proceedings of the 2nd international workshop on Computer Aided Verification (CAV'90)*, volume 531 of *Lecture Notes in Computer Science*, pages 233– 242. Springer-Verlag, 1991. (Cited on pages 3 and 12.)
- [27] Jean-Michel Couvreur. On-the-fly verification of temporal logic. In Jeannette M. Wing, Jim Woodcock, and Jim Davies, editors, *Proceedings of the World Congress on Formal Methods in the Development of Computing Systems (FM'99)*, volume 1708 of *Lecture Notes in Computer Science*, pages 253–271, Toulouse, France, September 1999. Springer-Verlag. ISBN 3-540-66587-0. (Cited on pages 24, 30, 33, 37, 155 and 158.)
- [28] Jean-Michel Couvreur. Un point de vue symbolique sur la logique temporelle linéaire. In Pierre Leroux, editor, Actes du Colloque LaCIM 2000, volume 27 of Publications du LaCIM, pages 131–140, Montréal, August 2000. Université du Québec à Montréal. (Cited on page 30.)
- [29] Jean-Michel Couvreur, Alexandre Duret-Lutz, and Denis Poitrenaud. On-the-fly emptiness checks for generalized Büchi automata. In Patrice Godefroid, editor, *Proceedings of the* 12th International SPIN Workshop on Model Checking of Software (SPIN'05), volume 3639 of Lecture Notes in Computer Science, pages 143–158. Springer, August 2005. (Cited on pages 30, 33 and 37.)

- [30] Christian Dax, Jochen Eisinger, and Felix Klaedtke. Mechanizing the powerset construction for restricted classes of ω-automata. In Kedar S. Namjoshi, Tomohiro Yoneda, Teruo Higashino, and Yoshio Okamura, editors, *Proceedings of the 5th International Symposium* on Automated Technology for Verification and Analysis (ATVA'07), volume 4762 of Lecture Notes in Computer Science. Springer, October 2007. (Cited on page 25.)
- [31] Michel Diaz. *Réseaux de Petri, Modèles fondamentaux*. Traité IC2, série Informatique et systèmes d'information. Hermes Science, June 2001. (Cited on page 13.)
- [32] Edsger Wybe Dijkstra. EWD 376: Finding the maximum strong components in a directed graph. http://www.cs.utexas.edu/users/EWD/ewd03xx/EWD376.PDF, May 1973. (Cited on pages 36 and 48.)
- [33] Edsger Wybe Dijkstra. Finding the maximal strong components in a directed graph. In A Discipline of Programming, chapter 25, pages 192–200. Prentice-Hall, 1976. (Cited on pages 36 and 48.)
- [34] Alexandre Duret-Lutz. LTL translation improvements in Spot. In Proceedings of the 5th International Workshop on Verification and Evaluation of Computer and Communication Systems (VECoS'11), Electronic Workshops in Computing, Tunis, Tunisia, September 2011. British Computer Society. URL http://ewic.bcs.org/category/15853. (Cited on page 54.)
- [35] Alexandre Duret-Lutz. LTL translation improvements in Spot 1.0. International Journal on Critical Computer-Based Systems, 5(1/2):31–54, March 2014. (Cited on page 39.)
- [36] Alexandre Duret-Lutz and Denis Poitrenaud. SPOT: an extensible model checking library using transition-based generalized Büchi automata. In *Proceedings of the 12th IEEE/ACM International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems (MASCOTS'04)*, pages 76–83. IEEE Computer Society Press, October 2004. (Cited on pages 5, 7, 39, 52, 53, 88, 108, 115, 122, 149 and 155.)
- [37] Alexandre Duret-Lutz, Kais Klai, Denis Poitrenaud, and Yann Thierry-Mieg. Combining explicit and symbolic approaches for better on-the-fly LTL model checking. Technical Report 1106.5700, arXiv, June 2011. URL http://arxiv.org/abs/1106.5700. Extended version of our ATVA'11 paper, presenting two new techniques instead of one. (Cited on pages 7, 134, 136, 137, 140, 141, 142, 143 and 158.)
- [38] Alexandre Duret-Lutz, Kais Klai, Denis Poitrenaud, and Yann Thierry-Mieg. Self-loop aggregation product — a new hybrid approach to on-the-fly LTL model checking. In *Proceedings of the 9th International Symposium on Automated Technology for Verification and Analysis (ATVA'11)*, volume 6996 of *Lecture Notes in Computer Science*, pages 336–350, Taipei, Taiwan, October 2011. Springer. (Cited on pages 7, 131, 133, 134, 135, 141, 143, 146 and 157.)
- [39] Kousha Etessami. Stutter-invariant languages, ω-automata, and temporal logic. In N. Halbwachs and D. Peled, editors, *Proceedings of the 11th International Conference on Computer*

Aided Verification (CAV'99), volume 1633 of Lecture Notes in Computer Science, pages 236–248. Springer-Verlag, 1999. (Cited on pages 3, 4, 18 and 40.)

- [40] Kousha Etessami and Gerard J. Holzmann. Optimizing Büchi automata. In C. Palamidessi, editor, *Proceedings of the 11th International Conference on Concurrency Theory (Concur'00)*, volume 1877 of *Lecture Notes in Computer Science*, pages 153–167, Pennsylvania, USA, 2000. Springer-Verlag. (Cited on pages 26, 28, 37 and 46.)
- [41] Kousha Etessami, Thomas Wilke, and Rebecca A. Schuller. Fair simulation relations, parity games, and state space reduction for Büchi automata. In Fernando Orejas, Paul G. Spirakis, and Jan van Leeuwen, editors, *Proceedings of the 28th international colloquium on Automata, Languages and Programming*, volume 2076 of *Lecture Notes in Computer Science*, pages 694–707, Crete, Greece, July 2001. Springer-Verlag. (Cited on page 28.)
- [42] Berndt Farwer. Automata logics, and infinite games, volume 2500 of Lecture Notes in Computer Science, chapter ω-automata, pages 3–21. Springer-Verlag, 2002. ISBN 3-540-003886. URL http://dl.acm.org/citation.cfm?id=938135.938137. (Cited on pages 2 and 19.)
- [43] Kathi Fisler, Ranan Fraer, Gila Kamhi, Moshe Y. Vardi, and Zijiang Yang. Is there a best symbolic cycle-detection algorithm? In *Proceedings of the 7th International Conference* on Tools and Algorithms for the Construction and Analysis of Systems, TACAS 2001, pages 420–434. Springer-Verlag, 2001. ISBN 3-540-41865-2. (Cited on pages 114, 115, 116 and 117.)
- [44] Carsten Fritz. Constructing Büchi automata from linear temporal logic using simulation relations for alternating Büchi automata. In Oscar H. Ibarra and Zhe Dang, editors, *Proceedings of the 8th International Conference on Implementation and Application of Automata* (CIAA'03), volume 2759 of *Lecture Notes in Computer Science*, pages 35–48, Santa Barbara, California, July 2003. Springer-Verlag. ISBN 3-540-40561-5. (Cited on pages 28 and 46.)
- [45] Paul Gastin and Denis Oddoux. Fast LTL to Büchi automata translation. In G. Berry, H. Comon, and A. Finkel, editors, *Proceedings of the 13th International Conference on Computer Aided Verification (CAV'01)*, volume 2102 of *Lecture Notes in Computer Science*, pages 53–65, Paris, France, 2001. Springer-Verlag. (Cited on page 30.)
- [46] Jaco Geldenhuys and Henri Hansen. Larger automata and less work for LTL model checking. In *Proceedings of the 13th International SPIN Workshop (SPIN'06)*, volume 3925 of *Lecture Notes in Computer Science*, pages 53–70. Springer, 2006. (Cited on pages 5, 6, 33, 40, 44, 45, 47, 48, 52, 54, 61, 64, 70, 75, 85, 155 and 156.)
- [47] Jaco Geldenhuys and Antti Valmari. Tarjan's algorithm makes on-the-fly LTL verification more efficient. In Kurt Jensen and Andreas Podelski, editors, *Proceedings of the 10th International Conference on Tools and Algorithms for the Construction and Analysis of Systems* (*TACAS'04*), volume 2988 of *Lecture Notes in Computer Science*, pages 205–219. Springer, 2004. ISBN 3-540-21299-X. (Cited on page 33.)

- [48] Rob Gerth, Doron Peled, Moshe Y. Vardi, and Pierre Wolper. Simple on-the-fly automatic verification of linear temporal logic. In *Proceedings of the 15th Workshop on Protocol Specification Testing and Verification (PSTV'95)*, pages 3–18, Warsaw, Poland, 1996. Chapman & Hall. URL http://citeseer.nj.nec.com/gerth95simple.html. (Cited on pages 3, 12, 20 and 30.)
- [49] Dimitra Giannakopoulou and Flavio Lerda. From states to transitions: Improving translation of LTL formulæ to Büchi automata. In D.A. Peled and M.Y. Vardi, editors, *Proceedings* of the 22nd IFIP WG 6.1 International Conference on Formal Techniques for Networked and Distributed Systems (FORTE'02), volume 2529 of Lecture Notes in Computer Science, pages 308–326, Houston, Texas, November 2002. (Cited on pages 5, 7, 23, 24, 30, 37, 88, 108 and 155.)
- [50] Patrice Godefroid. Partial-Order Methods for the Verification of Concurrent Systems: An Approach to the State-Explosion Problem, volume 1032 of Lecture Notes in Computer Science. Springer-Verlag, 1996. ISBN 3540607617. (Cited on pages 4, 37 and 161.)
- [51] Alexandre Hamez, Yann Thierry-Mieg, and Fabrice Kordon. Hierarchical set decision diagrams and automatic saturation. In *Proceedings of the 29th international conference on Applications and Theory of Petri Nets*, PETRI NETS '08, pages 211–230. Springer-Verlag, 2008. ISBN 978-3-540-68745-0. doi: 10.1007/978-3-540-68746-7\_16. (Cited on pages 119, 120 and 122.)
- [52] Henri Hansen, Wojciech Penczek, and Antti Valmari. Stuttering-insensitive automata for on-the-fly detection of livelock properties. In Rance Cleaveland and Hubert Garavel, editors, *Proceedings of the 7th International ERCIM Workshop in Formal Methods for Industrial Critical Systems (FMICS'02)*, volume 66(2) of *Electronic Notes in Theoretical Computer Science*, Málaga, Spain, July 2002. Elsevier. (Cited on pages 5, 40, 53 and 155.)
- [53] Monika Heiner, David Gilbert, and Robin Donaldson. Petri nets for systems and synthetic biology. In *Proceedings of the 8th International School on Formal Methods for the Design of Computer (SFM'08)*, volume 5016 of *Lecture Notes in Computer Science*, pages 215–264. Springer, 2008. (Cited on page 53.)
- [54] Gerard J. Holzmann. Software model checking. NATO Summer School, pages 309–355, Marktoberdorf, Germany, August 2000. IOS Press Computer and System Sciences. (Cited on page 1.)
- [55] Gerard J. Holzmann. The Spin Model Checker: Primer and Reference Manual. Addison-Wesley, 2003. ISBN 0-321-22862-6. (Cited on pages 4, 13, 16, 37 and 161.)
- [56] Jérôme Hugues, Yann Thierry-Mieg, Fabrice Kordon, Laurent Pautet, Soheib Barrir, and Thomas Vergnaud. On the formal verification of middleware behavioral properties. In Proceedings of the 9th International Workshop on Formal Methods for Industrial Critical Systems (FMICS'04), volume 133 of Electronic Notes in Theoretical Computer Science, pages 139–157. Elsevier Science Publishers, September 2004. (Cited on page 53.)

- [57] Yonit Kesten, Amir Pnueli, and Li on Raviv. Algorithmic verification of linear temporal logic specifications. In K.G. Larsen, S. Skyum, and G. Winskel, editors, *Proceedins of the 5th International Colloquium on Automata, Languages, and Programming (ICALP'98)*, volume 1443 of *Lecture Notes in Computer Science*, pages 1–16. Springer-Verlag, 1998. (Cited on pages 114, 116 and 117.)
- [58] Kais Klai and Denis Poitrenaud. MC-SOG: An LTL model checker based on symbolic observation graphs. In *Proceedings of the 29th International Conference on Application* and Theory of Petri Nets (ICATPN'08), volume 5062 of Lecture Notes in Computer Science, pages 288–306, Xi'an, China, June 2008. Springer. (Cited on pages 140, 141, 142 and 143.)
- [59] Fabrice Kordon, Alban Linard, Didier Buchs, Maximilien Colange, Sami Evangelista, Kai Lampka, Niels Lohmann, Emmanuel Paviot-Adet, Yann Thierry-Mieg, and Harro Wimmel. Report on the model checking contest at petri nets 2011. *T. Petri Nets and Other Models of Concurrency*, 6:169–196, 2012. (Cited on page 54.)
- [60] Leslie Lamport. What good is temporal logic? In *IFIP Congress*, pages 657–668, 1983. (Cited on page 4.)
- [61] Christof Löding. Efficient minimization of deterministic weak ω-automata. Information Processing Letters, 79(3):105–109, 2001. (Cited on page 46.)
- [62] Zohar Manna and Amir Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*. Springer-Verlag New York, Inc., New York, NY, USA, 1992. ISBN 0-387-97664-7. (Cited on page 16.)
- [63] Robert McNaughton. Testing and generating infinite sequences by a finite automaton. *Infor*mation and Control, 9(5):521–530, October 1966. (Cited on page 25.)
- [64] MoVe/LRDE. The Spot home page: http://spot.lip6.fr, 2014. (Cited on pages 4, 39, 40 and 52.)
- [65] David E. Muller. Infinite sequences and finite machines. In Proceedings of the 1963 Proceedings of the Fourth Annual Symposium on Switching Circuit Theory and Logical Design, SWCT '63, pages 3–16, Washington, DC, USA, 1963. IEEE Computer Society. doi: 10.1109/SWCT.1963.8. URL http://dx.doi.org/10.1109/SWCT.1963.8. (Cited on page 25.)
- [66] Denis Oddoux. Utilisation des automates alternants pour un model-checking efficace des logiques temporelles linéaires. PhD thesis, Universitée Paris 7, Paris, France, December 2003. (Cited on page 28.)
- [67] Radek Pelánek. BEEM: benchmarks for explicit model checkers. In *Proceedings of the* 14th international SPIN conference on Model checking software, Lecture Notes in Computer Science, pages 263–267. Springer, 2007. (Cited on pages 124 and 149.)
- [68] Radek Pelánek. Properties of state spaces and their applications. *International Journal on Software Tools for Technology Transfer (STTT)*, 10(5):443–454, 2008. (Cited on page 61.)

- [69] Doron Peled. Combining partial order reductions with on-the-fly model-checking. In Proceedings of the 6th International Conference on Computer Aided Verification (CAV'94), volume 818 of Lecture Notes in Computer Science, pages 377–390. Springer-Verlag, 1994. (Cited on pages 4, 37 and 161.)
- [70] Doron Peled and Thomas Wilke. Stutter-invariant temporal properties are expressible without the next-time operator. *Information Processing Letters*, 63(5):243–246, September 1995. (Cited on pages 4 and 18.)
- [71] Irfan Pyarali, Marina Spivak, Ron Cytron, and Douglas Clark Schmidt. Evaluating and optimizing thread pool strategies for RT-CORBA. In *Proceeding of the ACM SIGPLAN* workshop on Languages, compilers and tools for embedded systems (LCTES'00), pages 214– 222. ACM, 2000. (Cited on page 53.)
- [72] M. O. Rabin and D. Scott. Finite automata and their decision problems. *IBM J. Res. Dev.*, 3(2):114–125, April 1959. ISSN 0018-8646. doi: 10.1147/rd.32.0114. URL http://dx. doi.org/10.1147/rd.32.0114. (Cited on page 25.)
- [73] Etienne Renault, Alexandre Duret-Lutz, Fabrice Kordon, and Denis Poitrenaud. Three SCCbased emptiness checks for generalized Büchi automata. In Ken McMillan, Aart Middeldorp, and Andrei Voronkov, editors, *Proceedings of the 19th International Conference on Logic* for Programming, Artificial Intelligence, and Reasoning (LPAR'13), volume 8312 of Lecture Notes in Computer Science, pages 668–682. Springer, December 2013. (Cited on page 33.)
- [74] Kristin Y. Rozier and Moshe Y. Vardi. LTL satisfiability checking. In Proceedings of the 12th International SPIN Workshop on Model Checking of Software (SPIN'07), volume 4595 of Lecture Notes in Computer Science, pages 149–167. Springer, 2007. (Cited on page 54.)
- [75] Kristin Y. Rozier and Moshe Y. Vardi. A multi-encoding approach for LTL symbolic satisfiability checking. In *Proceedings of the 17th international conference on Formal methods* (*FM'11*), pages 417–431. Springer, 2011. (Cited on pages 114 and 115.)
- [76] S. Safra. On the complexity of omega -automata. In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, SFCS '88, pages 319–327, Washington, DC, USA, 1988. IEEE Computer Society. ISBN 0-8186-0877-3. doi: 10.1109/SFCS.1988.21948. URL http://dx.doi.org/10.1109/SFCS.1988.21948. (Cited on page 25.)
- [77] Stefan Schwoon and Javier Esparza. A note on on-the-fly verification algorithms. In Nicolas Halbwachs and Lenore Zuck, editors, *Proceedings of the 11th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'05)*, volume 3440 of *Lecture Notes in Computer Science*, Edinburgh, Scotland, April 2005. Springer. (Cited on page 33.)
- [78] Roberto Sebastiani and Stefano Tonetta. "more deterministic" vs. "smaller" Büchi automata for efficient LTL model checking. In G. Goos, J. Hartmanis, and J. van Leeuwen, editors, Proceedings of the 12th Advanced Research Working Conference on Correct Hardware Design and Verification Methods (CHARME'03), volume 2860 of Lecture Notes in Computer

*Science*, pages 126–140, L'Aquila, Italy, October 2003. Springer-Verlag. (Cited on pages 5, 30, 37 and 40.)

- [79] Roberto Sebastiani, Stefano Tonetta, and Moshe Y. Vardi. Symbolic systems, explicit properties: on hybrid approches for LTL symbolic model checking. In Kousha Etessami and Sriram K. Rajamani, editors, *Proceedings of 17th International Conference on Computer Aided Verification (CAV'05)*, volume 3576 of *Lecture Notes in Computer Science*, pages 350–363, Edinburgh, Scotland, UK, July 2005. Springer. (Cited on pages 7, 114, 115, 116, 122, 131, 133 and 157.)
- [80] F. Somenzi, K. Ravi, and R. Bloem. Analysis of symbolic SCC hull algorithms. In *Proc.* of *FMCAD*'02, volume 2517 of *LNCS*, pages 88–105. Springer, 2002. (Cited on pages 114 and 117.)
- [81] Fabio Somenzi and Roderick Bloem. Efficient Büchi automata for LTL formulæ. In Proceedings of the 12th International Conference on Computer Aided Verification (CAV'00), volume 1855 of Lecture Notes in Computer Science, pages 247–263, Chicago, Illinois, USA, 2000. Springer-Verlag. (Cited on pages 28 and 37.)
- [82] Robert Tarjan. Depth-first search and linear graph algorithms. In *Conference records of the 12th Annual IEEE Symposium on Switching and Automata Theory*, pages 114–121. IEEE, October 1971. Later republished as [83]. (Cited on page 48.)
- [83] Robert Tarjan. Depth-first search and linear graph algorithms. *SIAM Journal on Computing*, 1(2):146–160, 1972. (Cited on pages 48 and 171.)
- [84] Heikki Tauriainen. Automata and Linear Temporal Logic: Translation with Transition-based Acceptance. PhD thesis, Helsinki University of Technology, Espoo, Finland, September 2006. (Cited on pages 30 and 33.)
- [85] Yann Thierry-Mieg, Denis Poitrenaud, Alexandre Hamez, and Fabrice Kordon. Hierarchical set decision diagrams and regular models. In *Proceedings of the 15th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'09)*, volume 5505 of *Lecture Notes in Computer Science*, pages 1–15. Springer-Verlag, 2009. ISBN 978-3-642-00767-5. doi: 10.1007/978-3-642-00768-2\_1. (Cited on pages 7, 114, 115, 122 and 149.)
- [86] Antti Valmari. Stubborn sets for reduced state space generation. In Proceedings of the 10th International Conference on Applications and Theory of Petri Nets (ICATPN'91), volume 618 of Lecture Notes in Computer Science, pages 491–515, London, UK, 1991. Springer-Verlag. (Cited on pages 4, 37 and 161.)
- [87] Antti Valmari. The state explosion problem. In W. Reisig and G. Rozenberg, editors, *Lectures on Petri Nets 1: Basic Models*, volume 1491 of *Lecture Notes in Computer Science*, pages 429–528. Springer-Verlag, 1998. (Cited on pages 2, 12 and 155.)

- [88] Antti Valmari. Bisimilarity minimization in O(mlog n) time. In Proceedings of the 30th International Conference on the Applications and Theory of Petri Nets (ICATPN'09), volume 5606 of Lecture Notes in Computer Science, pages 123–142. Springer, 2009. ISBN 978-3-642-02423-8. URL http://dx.doi.org/10.1007/978-3-642-02424-5\_9. (Cited on pages 26 and 27.)
- [89] Moshe Y. Vardi. An automata-theoretic approach to linear temporal logic. In Faron Moller and Graham M. Birtwistle, editors, *Proceedings of the 8th Banff Higher Order Workshop* (*Banff'94*), volume 1043 of *Lecture Notes in Computer Science*, pages 238–266, Banff, Alberta, Canada, 1996. Springer-Verlag. ISBN 3-540-60915-6. (Cited on pages 2, 12 and 155.)
- [90] Moshe Y. Vardi. Automata-theoretic model checking revisited. In Proceedings of the 8th International Conference on Verification, Model Checking and Abstract Interpretation (VM-CAI'07), volume 4349 of Lecture Notes in Computer Science, Nice, France, January 2007. Springer. Invited paper. (Cited on pages 2 and 155.)
- [91] Ralf Wimmer, Marc Herbstritt, Holger Hermanns, Kelley Strampp, and Bernd Becker. Sigref: a symbolic bisimulation tool box. In *Proceedings of the 4th international conference on Automated Technology for Verification and Analysis*, ATVA'06, pages 477–492, Berlin, Heidelberg, 2006. Springer-Verlag. ISBN 3-540-47237-1, 978-3-540-47237-7. doi: 10.1007/11901914\_35. URL http://dx.doi.org/10.1007/11901914\_35. (Cited on pages 26 and 27.)
- [92] Pierre Wolper. Constructing automata from temporal logic formulas: A tutorial. In E. Brinksma, H. Hermanns, and J.-P. Katoen, editors, *Proceedings of the FMPA 2000 summer school*, volume 2090 of *Lecture Notes in Computer Science*, pages 261–277, Nijmegen, the Netherlands, July 2000. Springer-Verlag. (Cited on page 28.)
- [93] Pierre Wolper, Moshe Y. Vardi, and Aravinda Prasad Sistla. Reasoning about infinite computation paths. In *Proceedings of the 24th IEEE Symposium on Foundations of Computer Science (FOCS'83)*, pages 185–194. IEEE Computer Society Press, 1983. (Cited on page 28.)
## APPENDIX A

## Experimental Comparison of Explicit approaches using TGBA, TA and TGTA, with TGBA improved using simulation-reduction

Recently, several TGBA optimizations was added to SPOT, especially to apply the simulationreduction [3] to TGBA. The following tables and scatter plots show the impact of this optimization on the results of our experimental comparison (between TGBA, TA and TGTA) presented in Chapter 5.

Table A.1 and Table A.2 shows how for TGBA, TA and TGTA approaches deal with toy models and random formulas. We omit data for BA since they are always outperformed by TGBA. Table A.5 and Table A.6 show toy models against weak-fairness formulas.

Table A.3 and Table A.4 show the results of the two cases studies against random, weak-fairness, and dedicated formulas issued from the case studies (see Section 3.6.2).

These tables separate cases where formulas are verified from cases where they are violated. In the former (Tables A.1, A.5 and A.3), no counterexample are found and the full state-space had to be explored; in the latter (Tables A.2, A.6 and 5.4) the on-the-fly exploration of the state-space stopped as soon as the existence of a counterexample could be computed.

The column " $T_{\varphi}$ " shows the time (in  $\frac{1}{100^e}$  of seconds) spent constructing the property automata  $A_{\neg\varphi_i}$  from the formulas, this time includes the cost of the simulation-reduction of TGBA. This cost impacts all the other approaches because TGTA, BA and TA (through a BA) are constructed from TGBA.

Figure A.1 and Figure A.2 compare the number of visited transitions when running the emptiness check; plotting TGTA against respectively TGBA and TA.

			Verified properties (no counterexample)							
			Automaton			Full	product	Emptiness check		
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т
eterson5	TCDA	avg	9	327	21	2134328	8691523	2134328	8691523	13038
	TGBA	max	48	2250	1937	8634463	58408155	8634463	58408155	53292
	TA	avg	69	3444	29	2083706	8114637	3 578 356	13947438	22495
		max	531	31459	1960	7435026	29160195	14870052	58320390	89846
Ā	тота	avg	57	3880	30	1963093	7650369	1963093	7650369	12294
	IGIA	max	386	42415	1971	6635199	26179377	6635199	26179377	41328
	TODA	avg	9	273	2	1043179	6279946	1043179	6279946	1473
	TGBA	max	49	3550	22	2849784	21495868	2849784	21495868	4633
1 <u>g</u> 6		avg	54	2127	8	838315	4871788	1648030	9579646	2690
Rir		max	531	24608	154	2209660	12784696	4419320	25 569 392	6815
	TOTA	avg	44	2315	8	823663	4921456	823 663	4921456	1361
	IGIA	max	386	31457	180	2084839	12563232	2084839	12563232	3357
	TCDA	avg	7	188	4	1655712	10797107	1655712	10797107	2196
	IGBA	max	22	1651	286	8839019	73 557 865	8839019	73 557 865	15448
IS5	TA	avg	41	1360	8	1182686	9148797	1278218	9880397	2444
F		max	282	11609	295	5286393	41 375 144	6367350	50679714	11616
	TGTA	avg	33	1469	8	1138010	8829003	1138010	8829003	2085
		max	255	13384	301	5286393	41 375 144	5286393	41 375 144	9606
	TGBA	avg	6	121	1	2585110	20573569	2585110	20573569	3344
N		max	48	1994	11	17706640	149 153 508	17706640	149 153 508	22039
ban	<b>T</b> A	avg	38	1 308	3	1895866	17696106	1 895 866	17696106	3 5 6 2
an	IA	max	264	18440	80	10558520	102370471	10558520	102370471	19446
×	тста	avg	31	1 398	3	1802404	16858757	1802404	16858757	3244
	IUIA	max	220	22092	89	9345280	91 147 301	9345280	91 147 301	17228
	TGBA	avg	11	500	5	3970577	20450114	3970577	20450114	6638
0	TODA	max	73	5397	233	17947837	85727092	17947837	85727092	29036
101	ТА	avg	61	3 3 9 7	20	2012055	16929172	2012055	16929172	6067
Phi		max	404	51034	475	7557069	65 1 33 806	7 5 5 7 0 6 9	65 133 806	22797
	тста	avg	48	3767	23	1797978	15101356	1 797 978	15101356	5338
	IUIA	max	294	48954	444	6660936	54496416	6660936	54496416	19902
	тсра	avg	9	357	4	1847692	16517326	1847692	16517326	5022
N	IUDA	max	50	3284	61	7745523	77 343 545	7745523	77 343 545	23212
lin1	ТА	avg	65	3881	13	1468280	11959714	2829126	23057049	8151
Sob	IA	max	462	34568	137	6544384	52510720	13088768	105 021 440	35 548
	TGTA	avg	57	4518	14	1492636	12399229	1492636	12399229	4343
	IUIA	max	369	40112	133	6556672	54 543 361	6556672	54 543 361	18743

Table A.1: Comparison of the three approaches on toy examples with **random formulæ**, when counterexamples do not exist.

			Violated properties (a counterexample exists)								
			Automaton			Full p	product	Emptiness check			
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т	
n5	TGRA	avg	11	470	4	6875386	29707327	556882	2059495	3 3 0 8	
	TODA	max	47	2630	132	35221089	164860104	3266077	13122565	19494	
LSO	TA	avg	77	4616	15	7311696	28594154	531367	1847946	3 3 0 3	
ete	IA	max	688	62540	151	33610951	130319612	3034231	10910727	18973	
	тста	avg	58	4529	16	6724275	26406241	519338	1814387	3171	
	IUIA	max	377	39300	156	32 323 597	126519495	3 2 4 1 5 8 4	11733133	19719	
	тсра	avg	12	443	28	2366645	15175961	423 272	2286638	536	
	IGDA	max	51	2636	2038	7710378	56195862	2543904	14240224	3351	
1 <u>g</u> 6	<b>T</b> A	avg	96	5312	37	2181104	13018794	316826	1785138	546	
Ri	IA	max	435	36434	2069	7791264	47132088	1813390	10805447	3150	
	тста	avg	76	5389	38	1962430	12050193	304 280	1732216	476	
	IGIA	max	419	39766	2063	6516704	40390528	1709454	10552980	2790	
	тсра	avg	10	321	5	7030618	54559769	648095	2887241	649	
	IGDA	max	84	3034	236	28022123	252921593	9048245	63249152	13067	
IS5	TA	avg	68	2903	11	6336121	51452459	523 842	3433001	894	
E		max	396	20095	281	24916116	214008832	7336262	61 1 23 574	14439	
	TGTA	avg	55	2963	11	5726926	47023062	389 889	2490868	633	
		max	337	20595	301	23351734	205586222	6594559	54402238	12704	
	TGBA	avg	8	169	3	7 389 200	65265886	398 291	2413326	403	
N		max	25	1254	73	20380864	249816728	4154893	40027194	6514	
bar	<b>T</b> A	avg	50	1655	6	6099926	58043214	279149	2293392	467	
an	IA	max	271	15866	81	20138608	208914286	2877959	27 297 928	5 5 8 5	
×.	тста	avg	41	1754	7	5664045	54798536	255545	2095285	410	
	IUIA	max	228	18303	80	15816808	170535764	2877959	27 297 927	5294	
	тсра	avg	9	389	2	14648744	127 225 739	746849	3135235	1 1 0 8	
	TODA	max	52	3832	20	67346113	944 250 350	9032250	43945324	14928	
<b>I</b> 01	ТА	avg	71	3877	11	12807603	115165982	514894	3437571	1324	
Phi	IA	max	272	34438	103	53 294 292	530498041	7342016	64470778	22435	
.	тста	avg	58	4118	12	11611005	105837118	434875	2948356	1098	
	IUIA	max	268	35 598	112	55 198 542	558520250	3 6 2 2 3 4 5	30490183	11529	
	TGRA	avg	14	543	16	4072114	37723236	514761	5099096	1 506	
N	TODA	max	50	3832	529	12193720	166810859	10485865	153206175	41271	
lii1	T۸	avg	122	7708	28	3606864	29540049	473 503	3870325	1446	
Sob	IA	max	799	55662	548	8984576	81 599 488	7108784	66618546	22451	
	TGTA	avg	98	8057	29	3 390 389	28308511	489 600	4058013	1406	
	IUIA	max	485	64827	561	9922560	91222016	8046768	76224690	24848	

Table A.2: Comparison of the three approaches on toy examples with **random formulæ**, when **counterexamples exist.** 

				Verified properties (no counterexample)								
				1	Automato	n	Full	product	Emptiness check			
				st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т	
		TGBA	avg	10	429	20	85 195	207750	85 195	207750	425	
			max	50	3284	1098	238118	1127918	238118	1127918	1248	
	Ð	ТЛ	avg	67	3787	29	78382	168410	128674	276738	655	
	2	IA	max	435	34568	1137	202874	436381	405748	872762	2017	
		тста	avg	55	4045	31	81656	176671	81 656	176671	413	
3/2		IUIA	max	323	30547	1160	221779	485262	221779	485 262	1090	
3 3/		тсра	avg	3	42	1	73879	168772	73879	168772	386	
RI		IUDA	max	18	292	5	181518	619068	181518	619068	968	
<b>JyC</b>	air	TA	avg	67	850	1	95006	204347	95623	205 667	508	
$\mathbf{P}_{0}$	8		max	353	5953	11	287185	630231	287 185	630231	1 499	
		TGTA	avg	29	420	1	70158	150547	70158	150547	370	
			max	99	1700	4	132396	286698	132396	286698	698	
	$\Phi_1$	TGBA	-	4	208	2	221916	496626	221916	496626	1070	
		TA	-	68	5677	3	219288	478950	438576	957900	2088	
		TGTA	_	67	7556	5	221952	488136	221952	488136	1052	
		TGBA	avg	9	246	2	2598503	22108346	2598503	22108346	4833	
			max	64	3551	16	11272599	120406938	11272599	120406938	22707	
	Ð		avg	42	1651	8	1608896	18824425	1608896	18824425	4583	
	2	IA	max	336	18971	157	8036021	100718455	8036021	100718455	23124	
		тста	avg	35	1876	9	1519337	17816286	1519337	17816286	4227	
		IGIA	max	317	25595	211	8036021	100718455	8036021	100718455	21791	
K 8		тера	avg	4	31	1	3461837	28114156	3461837	28114156	5 572	
<b>AP</b>		IUDA	max	21	160	3	10871182	104579823	10871182	104 579 823	20539	
N	air	T۸	avg	42	386	1	1961055	22640584	1961055	22640584	5129	
	8	IA	max	167	1902	4	6110748	75624744	6110748	75624744	17114	
		тста	avg	19	203	1	1777295	20615811	1777295	20615811	4550	
		IGIA	max	80	1052	3	9523917	118607814	9523917	118607814	26652	
		TGBA	-	5	152	0	46493	325442	46493	325 442	51	
	$ \Phi $	TA	-	7	210	2	33376	291602	33 376	291602	58	
		TGTA	-	6	343	3	33 376	291602	33 376	291602	55	

Table A.3: Comparison of the three approaches for the **case studies when counterexamples do not exist.** 

						V	olated properties (a counterexample exists)				
				Automaton			Full	product	Emptiness check		
				st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т
		TGBA	avg	10	503	6	113778	281 259	57622	148161	289
			max	56	4470	224	772777	1787484	279266	851459	1406
	Ð	Т۸	avg	60	4350	24	118209	256484	56406	126308	289
	2	IA	max	472	46228	368	1005017	2159320	208726	487016	1037
		TGTA	avg	50	5072	27	114 499	249711	$\mathbf{47468}$	105349	242
312			max	310	55353	520	758720	1627829	105707	240336	532
33/		TGBA	avg	3	41	1	78931	177210	63721	144 633	333
R			max	13	203	6	187020	533673	169655	476766	885
NC I	Fair	TA	avg	71	948	1	180036	386413	107 127	234966	566
Po	$\mathbb{R}$		max	204	3205	6	483707	1034329	245740	539608	1274
		TGTA	avg	30	452	1	78297	167718	$\boldsymbol{63504}$	138717	334
			max	102	1719	5	186891	399 306	169407	366 693	876
		тера	avg	9	321	3	20240822	245726563	730667	3 5 3 4 1 2 7	894
		IUDA	max	41	3012	68	77833929	1306979302	6547219	51 103 385	12377
	Ð	T۸	avg	62	3294	12	17057208	219712950	425966	2967248	872
	R	IA	max	314	26248	163	81988164	1146895628	4859218	31 565 707	8378
		тста	avg	50	3376	14	15139419	196274725	424117	2957187	829
		IUIA	max	229	29261	195	47 574 695	644775125	4859218	31 565 707	8 1 9 2
<b>X</b>		тсра	avg	4	52	1	15189697	191 231 582	1049014	6622077	1405
AP		TODA	max	13	221	3	53 57 5 3 1 3	764671391	7387837	64113702	12094
M/	Fair	Т۸	avg	76	986	2	22798869	291 635 140	638187	5996662	1444
	$\mathbb{R}$	IA	max	245	3322	5	64763487	858891872	5203746	59642224	13725
		тста	avg	34	503	1	12475839	160851874	545160	5000866	1183
		IUIA	max	87	1536	4	47723696	662295506	4859666	47087918	10019

Table A.4: Comparison of the three approaches for the case studies when counterexamples exist.

			Verified properties (no counterexample)							
			Automaton		Full product		Eı			
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т
eterson5	TGBA	avg	3	32	1	2334992	9214963	2334992	9214963	14219
		max	9	126	4	6197238	28678473	6197238	28678473	37883
	TA	avg	44	453	1	2594244	10057164	2915489	11303640	18390
		max	173	2360	4	10632587	41615545	10632587	41615545	66358
P	TGTA	avg	22	291	1	2205610	8569393	2205610	8569393	13744
		max	72	1103	5	5458289	21093007	5458289	21093007	33731
	тара	avg	4	38	1	885670	5085387	885670	5085387	1 3 0 9
Ì	IGBA	max	16	118	4	3095400	16216920	3095400	16216920	3934
1g6	T 4	avg	49	516	1	919252	5187969	1452203	8 196 805	2454
Rii		max	149	2382	5	2120000	12695504	3 565 136	20381440	5234
	тста	avg	26	347	1	725434	4184918	725434	4184918	1256
	IUIA	max	73	1 208	5	1300256	7556256	1 300 256	7556256	2432
	тера	avg	4	33	1	2622090	18893040	2622090	18893040	3733
		max	11	177	2	8685055	78815271	8685055	78815271	14741
IS5	T.4	avg	46	483	1	2045253	15834452	2045253	15834452	3947
F	IA	max	167	4256	3	7504266	60774774	7 504 266	60774774	14895
	TGTA	avg	22	265	1	1769028	13854565	1769028	13854565	3291
		max	87	1430	4	6949551	56088708	6949551	56088708	14382
	TGBA	avg	3	27	1	2679040	22247651	2679040	22247651	3464
Ŋ		max	21	168	3	9183294	87988390	9183294	87988390	12204
ban	T 4	avg	33	277	1	1887122	17126840	1887122	17126840	3 5 0 6
an	IA	max	136	1901	3	8429960	80660846	8429960	80660846	13034
X	тста	avg	17	177	1	1711068	15682969	1711068	15682969	3072
	IUIA	max	62	912	3	6965000	65796520	6965000	65796520	11364
	тсра	avg	3	28	1	3980003	23939240	3980003	23939240	7 599
0	TODA	max	11	154	2	12974557	112561242	12974557	112561242	31701
<b>lo1</b>	ТА	avg	45	474	1	3072870	26224843	3072870	26224843	9283
Phi		max	289	8154	6	10987384	94141317	10987384	94141317	31547
	тста	avg	20	239	1	2291618	19611203	2 291 618	19611203	6750
		max	116	3 3 0 3	3	9571804	85885107	9571804	85885107	28044
	тсра	avg	3	35	1	1362987	11158101	1 362 987	11158101	3736
N.		max	10	150	4	3059706	28950456	3059706	28950456	8801
lin1	ТА	avg	46	489	1	1235467	9824773	1453687	11553559	4482
Rob	IA	max	185	2558	5	2684928	21 341 184	4696064	37017600	13856
	TGTA	avg	24	323	1	1126156	9031138	1126156	9031138	3450
		max	92	1334	5	2211840	17928192	2211840	17928192	6866

Table A.5: Comparison of the three approaches on toy examples with **weak-fairness formulæ**, **when counterexamples do not exist.** 

			Violated properties (a counterexample exists)								
			Automaton			Full p	product	Emptiness check			
			st.	tr.	$T_{\phi}$	st.	tr.	st.	tr.	Т	
eterson5	TGBA	avg	4	48	1	2993058	11761446	471940	1504950	2793	
	TODA	max	18	292	8	6711952	32894473	2933653	10800747	17392	
	TA TGTA	avg	81	1156	1	9841905	38 327 384	567949	1843300	3488	
		max	352	6001	11	15973169	62098442	4444487	17106565	27499	
<b>P</b>		avg	34	534	1	3055186	11931357	469284	1501463	2828	
		max	117	2108	9	6405145	24914480	2898184	10706849	17570	
	TGBA	avg	5	68	1	1429086	8845224	254444	1 395 698	364	
	TGBA	max	28	482	7	3 0 3 2 6 6 4	21798322	1416696	7920367	2086	
1 <u>g</u> 6	TA	avg	107	1477	2	2871519	16768200	309870	1710585	526	
Ri	IA	max	322	6456	12	5107450	30709648	2397874	13900223	4221	
	тста	avg	44	697	2	1193721	7025538	215638	1222837	364	
	IUIA	max	178	3687	9	2 501 240	14944456	1 1 69 3 57	6706261	2103	
	тсра	avg	4	41	1	6039757	48073683	684361	3016021	660	
	IUDA	max	14	200	4	20600805	163 375 160	6676584	53451994	9502	
1S5	TA	avg	62	719	1	8698172	69734251	494830	3078804	799	
E		max	245	3136	5	25112720	206558998	6425987	51962113	11444	
	TGTA	avg	29	393	1	4733810	38333115	403379	2445137	616	
		max	94	1439	4	13 122 252	109888666	5062934	41766785	8867	
	TGBA	avg	4	42	1	5411516	48648241	615574	3554349	606	
N		max	15	161	3	17194688	155902603	7950633	80041971	10547	
bar	TA	avg	60	701	1	6955248	65524457	319238	2608392	532	
an	IA	max	208	3163	4	19860365	194830038	5 3 2 1 5 9 8	52675827	9208	
12	тста	avg	28	380	1	4023776	38395802	287093	2354941	457	
	IUIA	max	85	1228	4	10822336	106435273	5049942	50190351	8374	
	TGRA	avg	4	45	1	8488406	70073588	624958	2345810	837	
	TODA	max	12	164	4	21620277	175367424	4214484	21334294	6487	
<b>1</b> 0	ТА	avg	71	858	2	14875690	131208038	573945	3466645	1315	
Phi	17	max	248	3614	6	42 242 159	375 507 068	6494793	50994037	14508	
.	тста	avg	33	456	1	6888581	61271105	394840	2161711	837	
	IUIA	max	94	1431	4	16961885	150416553	2787217	21746406	6346	
	тсра	avg	5	76	2	2822635	25680929	439215	3708242	1215	
N	IUDA	max	28	482	7	8204253	114 291 493	3244925	30724614	9443	
lii1	Т^	avg	112	1610	3	4439903	35 584 823	437337	3439416	1347	
Sob	IA	max	361	7372	12	16381952	141807616	3 174 178	25462895	10666	
	тста	avg	46	761	2	2333623	18914483	364484	2908781	1109	
	IGIA	max	178	3687	9	6628864	56228864	2678258	21824221	8646	

Table A.6: Comparison of the three approaches on toy examples with **weak-fairness formulæ**, **when counterexamples exist.** 



(a) TGTA against TGBA approaches for verified properties



(b) TGTA against TGBA approaches for violated properties

Figure A.1: Performance of TGTA against TGBA, with TGBA improved using simulation-reduction.



(a) TGTA against TA approaches for verified properties



(b) TGTA against TA approaches for violated properties

Figure A.2: Performance of TGTA against TA (with TGBA improved using simulation-reduction).