

Simulations in Rank-Based Büchi Automata Complementation

Yu-Fang Chen¹, Vojtěch Havlena², and Ondřej Lengál²

¹ Academia Sinica, Taiwan

² FIT, IT4I Centre of Excellence, Brno University of Technology, Czech Republic

Abstract. Complementation of Büchi automata is an essential technique used in some approaches for termination analysis of programs. The long search for an optimal complementation construction climaxed with the work of Schewe, who proposed a worst-case optimal rank-based procedure that generates complements of a size matching the theoretical lower bound of $(0.76n)^n$, modulo a polynomial factor of $\mathcal{O}(n^2)$. Although worst-case optimal, the procedure in many cases produces automata that are unnecessarily large. In this paper, we propose several ways of how to use the direct and delayed simulation relations to reduce the size of the automaton obtained in the rank-based complementation procedure. Our techniques are based on either (i) ignoring macrostates that cannot be used for accepting a word in the complement or (ii) saturating macrostates with simulation-smaller states, in order to decrease their total number. We experimentally showed that our techniques can indeed considerably decrease the size of the output of the complementation.

1 Introduction

Büchi automata (BA) complementation is a fundamental problem in program analysis and formal verification, from both theoretical and practical angles. It is, for instance, a critical step in some approaches for termination analysis, which is an essential part of establishing total correctness of programs [14,19,9]. Moreover, BA complementation is used as a component of decision procedures of some logics for reasoning about programs, such as S1S capturing a decidable fragment of second-order arithmetic [6] or the temporal logics ETL and QPTL [35].

The study of the BA complementation problem can be traced back to 1962, when Büchi introduced his automaton model in the seminal paper [6] in the context of a decision procedure for the S1S fragment of second-order arithmetic. In the paper, a doubly exponential complementation algorithm based on the infinite Ramsey theorem is proposed. In 1988, Safra [32] introduced a complementation procedure with an $n^{\mathcal{O}(n)}$ upper bound and, in the same year, Michel [28] established an $n!$ lower bound. From the traditional theoretical point of view, the problem was already solved, since exponents in the two bounds matched under the \mathcal{O} notation (recall that $n!$ is approximately $(n/e)^n$). From a more practical point of view, a linear factor in an exponent has a significant impact on real-world applications. It was established that the upper bound of Safra's construction is 2^{2n} , so the hunt for an optimal algorithm continued [38]. A series

of research efforts participated in narrowing the gap [24,15,39,23,41]. The long journey climaxed with the result of Schewe [33], who proposed an optimal rank-based procedure that generates complements of a size matching the theoretical lower bound of $(0.76n)^n$ found by Yan [41], modulo a polynomial factor of $\mathcal{O}(n^2)$.

Although the algorithm of Schewe is worst-case optimal, it often generates unnecessarily large complements. The standard approach to alleviate this problem is to decrease the size of the input BA before the complementation starts. Since minimization of (nondeterministic) BAs is a PSPACE-complete problem, more lightweight reduction methods are necessary. The most prevalent approaches are those based on various notions of *simulation-based reduction*, such as reductions based on *direct simulation* [7,36], a richer *delayed simulation* [12], or their *multi-pebble* variants [13]. These approaches first compute a simulation relation over the input BA—which can be done with the time complexity $\mathcal{O}(mn)$ [20,22,30,31,8] and $\mathcal{O}(mn^3)$ [12] for direct and delayed simulation respectively, with the number of states n and transitions m —and then construct a *quotient* BA by merging simulation-equivalent states, while preserving the language of the input BA. The other approach is a reduction based on *fair simulation* [18]. The fair simulation cannot, however, be used for quotienting, but still it can be used for merging certain states and removing transitions. The reduced BA is used as the input of the complementation, which often significantly reduces the size of the result.

In this paper, we propose several ways of how to exploit the direct and delayed simulations in BA complementation even further to obtain smaller complements and shorter running times. We focus, in particular, on the optimal *rank-based* complementation procedure of Schewe [33]. Essentially, the rank-based construction is an extension of traditional subset construction for determinizing finite automata, with some additional information kept in each macrostate (a state in the complemented BA) to track the acceptance condition of all runs of the input automaton on a given word. In particular, it stores the *rank* of each state in a macrostate, which, informally, measures the distance to the last accepting state on the corresponding run in the input BA. The main contributions of this paper are the following optimisations of rank-based complementation for BAs, for an input BA \mathcal{A} and the output of the rank-based complementation algorithm \mathcal{B} .

1. *Purging*: We use simulation relations over \mathcal{A} to remove some useless macrostates during the construction of \mathcal{B} . In particular, if a state p is simulated by q in \mathcal{A} , this puts a restriction on the relation between the ranks of runs from p and from q . As a consequence, macrostates that assign ranks violating this restriction can be purged from \mathcal{B} .
2. *Saturation*: We saturate macrostates with states that are simulated by the macrostate; this can reduce the total number of states of \mathcal{B} because two or more macrostates can be mapped to a single saturated macrostate. This is inspired by the technique of Glabbeek and Ploeger that uses *closures* in finite automata determinization [17].

The proposed optimizations are orthogonal to simulation-based size reduction mentioned above. Since the quotienting methods are based on taking only the

symmetric fragment of the simulation, i.e., they merge states that simulate *each other*, after the quotienting, there might still be many pairs where the simulation holds in only one way, and can therefore be exploited by our techniques. Since the considered notions of simulation-based quotienting preserve the respective simulations, our techniques can be used to optimize the complementation *at no additional cost*. Our experimental evaluation of the optimizations showed that in many cases, they indeed significantly reduce the size of the complemented BA.

2 Preliminaries

We fix a finite nonempty alphabet Σ and the first infinite ordinal $\omega = \{0, 1, \dots\}$. For $n \in \omega$, by $[n]$ we denote the set $\{0, \dots, n\}$. An (infinite) word α is represented as a function $\alpha : \omega \rightarrow \Sigma$ where the i -th symbol is denoted as α_i . A finite word w of length $n + 1$ is represented as a function $w : [n] \rightarrow \Sigma$. The finite word of length 0 is denoted as ϵ . We abuse notation and sometimes also represent α as an infinite sequence $\alpha = \alpha_0\alpha_1\dots$ and w as a finite sequence $w = w_0\dots w_{n-1}$. The suffix $\alpha_i\alpha_{i+1}\dots$ of α is denoted by $\alpha_{i:\omega}$. We use Σ^ω to denote the set of all infinite words over Σ and Σ^* to denote the set of all finite words. For $L \subseteq \Sigma^*$ we define $L^* = \{u \in \Sigma^* \mid u = w_1\dots w_n \wedge \forall 1 \leq i \leq n : w_i \in L\}$ and $L^\omega = \{\alpha \in \Sigma^\omega \mid \alpha = w_1w_2\dots \wedge \forall i \geq 1 : w_i \in L\}$ (note that $\{\epsilon\}^\omega = \emptyset$). Given $L_1, L_2 \subseteq \Sigma^*$, we use L_1L_2 to denote the set $\{w_1w_2 \mid w_1 \in L_1, w_2 \in L_2\}$.

A (nondeterministic) *Büchi automaton* (BA) over Σ is a quadruple $\mathcal{A} = (Q, \delta, I, F)$ where Q is a finite set of *states*, δ is a *transition function* $\delta : Q \times \Sigma \rightarrow 2^Q$, and $I, F \subseteq Q$ are the sets of *initial* and *accepting* states respectively. We sometimes treat δ as a set of transitions $p \xrightarrow{a} q$, for instance, we use $p \xrightarrow{a} q \in \delta$ to denote that $q \in \delta(p, a)$. Moreover, we extend δ to sets of states $P \subseteq Q$ as $\delta(P, a) = \bigcup_{p \in P} \delta(p, a)$. A *run* of \mathcal{A} from $q \in Q$ on an input word α is an infinite sequence $\rho : \omega \rightarrow Q$ that starts in q and respects δ , i.e., $\rho_0 = q$ and $\forall i \geq 0 : \rho_i \xrightarrow{\alpha_i} \rho_{i+1} \in \delta$. We say that ρ is *accepting* iff it contains infinitely many occurrences of some accepting state, i.e., $\exists q_f \in F : |\{i \in \omega \mid \rho_i = q_f\}| = \omega$. A word α is *accepted* by \mathcal{A} from a state $q \in Q$ if there is an accepting run ρ of \mathcal{A} from q , i.e., $\rho_0 = q$. The set $\mathcal{L}_{\mathcal{A}}(q) = \{\alpha \in \Sigma^\omega \mid \mathcal{A} \text{ accepts } \alpha \text{ from } q\}$ is called the *language* of q (in \mathcal{A}). Given a set of states $R \subseteq Q$, we define the language of R as $\mathcal{L}_{\mathcal{A}}(R) = \bigcup_{q \in R} \mathcal{L}_{\mathcal{A}}(q)$ and the language of \mathcal{A} as $\mathcal{L}(\mathcal{A}) = \mathcal{L}_{\mathcal{A}}(I)$. For a pair of states p and q in \mathcal{A} , we use $p \subseteq_{\mathcal{L}} q$ to denote $\mathcal{L}_{\mathcal{A}}(p) \subseteq \mathcal{L}_{\mathcal{A}}(q)$.

Without loss of generality, in this paper, we assume \mathcal{A} to be complete, i.e., for every state q and symbol a , it holds that $\delta(q, a) \neq \emptyset$. A *trace* over a word α is an infinite sequence $\pi = q_0 \xrightarrow{\alpha_0} q_1 \xrightarrow{\alpha_1} \dots$ such that $\rho = q_0q_1\dots$ is a run of \mathcal{A} over α from q_0 . We say π is *fair* if it contains infinitely many accepting states. Moreover, we use $p \xrightarrow{w} q$ for $w \in \Sigma^*$ to denote that q is reachable from p over the word w ; if a path from p to q over w contains an accepting state, we can write $p \xrightarrow[F]{w} q$. In this paper, we fix a complete BA $\mathcal{A} = (Q, \delta, I, F)$.

2.1 Simulations

We introduce simulation relations between states of a BA \mathcal{A} using the game semantics in a similar manner as in the extensive study of Clemente and Mayr [26].

In particular, in a *simulation game* between two players (called Spoiler and Duplicator) in \mathcal{A} from a pair of states (p_0, r_0) , for any (infinite) trace over a word α that Spoiler takes starting from p_0 , Duplicator tries to mimic the trace starting from r_0 . On the other hand, Spoiler tries to find a trace that Duplicator cannot mimic. The game starts in the configuration (p_0, r_0) and every i -th round proceeds by, first, Spoiler choosing a transition $p_i \xrightarrow{\alpha_i} p_{i+1}$ and, second, Duplicator mimicking Spoiler by choosing a matching transition $r_i \xrightarrow{\alpha_i} r_{i+1}$ over the same symbol α_i . The next game configuration is (p_{i+1}, r_{i+1}) . Suppose that $\pi_p = p_0 \xrightarrow{\alpha_0} p_1 \xrightarrow{\alpha_1} \dots$ and $\pi_r = r_0 \xrightarrow{\alpha_0} r_1 \xrightarrow{\alpha_1} \dots$ are the two (infinite) traces constructed during the game. Duplicator *wins* the simulation game if $\mathcal{C}^x(\pi_p, \pi_r)$ holds, where $\mathcal{C}^x(\pi_p, \pi_r)$ is a condition that depends on the particular simulation. In the current paper, we consider the following simulation relations:

- **direct** [11]: $\mathcal{C}^{di}(\pi_p, \pi_r) \stackrel{\text{def}}{\iff} \forall i : p_i \in F \Rightarrow r_i \in F,$
- **delayed** [12]: $\mathcal{C}^{de}(\pi_p, \pi_r) \stackrel{\text{def}}{\iff} \forall i : p_i \in F \Rightarrow \exists k \geq i : r_k \in F, \text{ and}$
- **fair** [21]: $\mathcal{C}^f(\pi_p, \pi_r) \stackrel{\text{def}}{\iff} \text{if } \pi_p \text{ is fair, then } \pi_r \text{ is fair.}$

A maximal x -simulation relation $\preceq_x \subseteq Q \times Q$, for $x \in \{di, de, f\}$, is defined such that $p \preceq_x r$ iff Duplicator has a winning strategy in the simulation game with the winning condition \mathcal{C}^x starting from (p, r) . Formally, we define a strategy to be a (total) mapping $\sigma : Q \times (Q \times \Sigma \times Q) \rightarrow Q$ such that $\sigma(r, p \xrightarrow{a} p') \in \delta(r, a)$, i.e., if Duplicator is in state r and Spoiler selects a transition $p \xrightarrow{a} p'$, the strategy picks a state r' such that $r \xrightarrow{a} r' \in \delta$ (and because \mathcal{A} is complete, such a transition always exists). Note that Duplicator cannot look ahead at Spoiler's future moves. We use σ_x to denote any winning strategy of Duplicator in the \mathcal{C}^x simulation game. Let σ_x and σ'_x be a pair of winning strategies in the \mathcal{C}^x simulation game. We say that σ_x is *dominated* by σ'_x if for all states p and all transitions $q \xrightarrow{a} q'$ it holds that $\sigma_x(p, q \xrightarrow{a} q') \preceq_x \sigma'_x(p, q \xrightarrow{a} q')$, and that σ_x is *strictly dominated* by σ'_x if σ_x is dominated by σ'_x and σ_x does not dominate σ'_x . A strategy is *dominating* if it is not strictly dominated by any other strategy. Strategies are also lifted to traces as follows: let π_p be as above, then $\sigma(r_0, \pi_p) = r_0 \xrightarrow{\alpha_0} r_1 \xrightarrow{\alpha_1} \dots$ where for all $i \geq 0$ it holds that $\sigma(r_i, p_i \xrightarrow{\alpha_i} p_{i+1}) = r_{i+1}$. The considered simulation relations form the following hierarchy: $\preceq_{di} \subseteq \preceq_{de} \subseteq \preceq_f \subseteq \preceq_{\mathcal{L}}$. Note that every maximal simulation relation is a preorder, i.e., reflexive and transitive.

2.2 Run DAGs

In this section, we recall the terminology from [33] (which is a minor modification of the terminology from [24]). We fix the definition of the *run DAG* of \mathcal{A} over a word α to be a DAG (directed acyclic graph) $\mathcal{G}_\alpha = (V, E)$ of vertices V and edges E where

- $V \subseteq Q \times \omega$ s.t. $(q, i) \in V$ iff there is a run ρ of \mathcal{A} over α with $\rho_i = q$,
- $E \subseteq V \times V$ s.t. $((q, i), (q', i')) \in E$ iff $i' = i + 1$ and $q' \in \delta(q, \alpha_i)$.

Given \mathcal{G}_α as above, we will write $(p, i) \in \mathcal{G}_\alpha$ to denote that $(p, i) \in V$. We call (p, i) *accepting* if p is an accepting state. \mathcal{G}_α is *rejecting* if it contains no path with infinitely many accepting vertices. A vertex $(p, i) \in \mathcal{G}_\alpha$ is *finite* if the set of

vertices reachable from (p, i) is finite, *infinite* if it is not finite, and *endangered* if (p, i) cannot reach an accepting vertex.

We assign ranks to vertices of run DAGs as follows: Let $\mathcal{G}_\alpha^0 = \mathcal{G}_\alpha$ and $j = 0$. Repeat the following steps until the fixpoint or for at most $2n + 1$ steps, where n is the number of states of \mathcal{A} .

- Set $\text{rank}_\alpha(p, i) := j$ for all finite vertices (p, i) of \mathcal{G}_α^j and let \mathcal{G}_α^{j+1} be \mathcal{G}_α^j minus the vertices with the rank j .
- Set $\text{rank}_\alpha(p, i) := j + 1$ for all endangered vertices (p, i) of \mathcal{G}_α^{j+1} and let \mathcal{G}_α^{j+2} be \mathcal{G}_α^{j+1} minus the vertices with the rank $j + 1$.
- Set $j := j + 2$.

For all vertices v that have not been assigned a rank yet, we assign $\text{rank}_\alpha(v) := \omega$. (Note that since \mathcal{A} is complete, then $\mathcal{G}_\alpha^1 = \mathcal{G}_\alpha^0$.)

Lemma 1. *If $\alpha \notin \mathcal{L}(\mathcal{A})$, then $0 \leq \text{rank}_\alpha(v) \leq 2n$ for all $v \in \mathcal{G}_\alpha$. Moreover, if $\alpha \in \mathcal{L}(\mathcal{A})$, then there is a vertex $(p, 0) \in \mathcal{G}_\alpha$ s.t. $\text{rank}_\alpha(p, 0) = \omega$.*

Proof. Follows from Corollary 3.3 in [24]. □

3 Complementing Büchi Automata

We use as the starting point the complementation procedure of Schewe [33, Section 3.1], which we denote as COMP_S (the ‘S’ stands for ‘Schewe’). The procedure works with the notion of level rankings. Given $n = |Q|$, a (*level*) *ranking* is a function $f : Q \rightarrow [2n]$ such that $\{f(q_f) \mid q_f \in F\} \subseteq \{0, 2, \dots, 2n\}$, i.e., f assigns even ranks to accepting states of \mathcal{A} .³ For a ranking f , the *rank* of f is defined as $\text{rank}(f) = \max\{f(q) \mid q \in Q\}$. For a set of states $S \subseteq Q$, we call f to be *S-tight* if (i) it has an odd rank r , (ii) $\{f(s) \mid s \in S\} \supseteq \{1, 3, \dots, r\}$, and (iii) $\{f(q) \mid q \notin S\} = \{0\}$. A ranking is *tight* if it is Q -tight; we use \mathcal{T} to denote the set of all tight rankings. For a pair of rankings f and f' , a set $S \subseteq Q$, and a symbol $a \in \Sigma$, we use $f' \leq_a^S f$ iff for every $q \in S$ and $q' \in \delta(q, a)$ it holds that $f'(q') \leq f(q)$.

The COMP_S procedure constructs the BA $\mathcal{B}_S = (Q', \delta', I', F')$ whose components are defined as follows:

- $Q' = Q_1 \cup Q_2$ where
 - $Q_1 = 2^Q$ and
 - $Q_2 = \{(S, O, f, i) \in 2^Q \times 2^Q \times \mathcal{T} \times \{0, 2, \dots, 2n - 2\} \mid f \text{ is } S\text{-tight}, O \subseteq S \cap f^{-1}(i)\}$,
- $I' = \{I\}$,
- $\delta' = \delta_1 \cup \delta_2 \cup \delta_3$ where
 - $\delta_1 : Q_1 \times \Sigma \rightarrow 2^{Q_1}$ such that $\delta_1(S, a) = \{\delta(S, a)\}$,

³ Note that our basic definitions slightly differs from the ones in Section 2.3 of [33]. This is because of a typo in [33]; indeed, if the procedure from [33] is implemented as is, the output does not accept the complement (there might be a macrostate (S, O, f) where S contains accepting states and O is empty, and, therefore, the whole macrostate is accepting, which is wrong).

- $\delta_2 : Q_1 \times \Sigma \rightarrow 2^{Q_2}$ such that $\delta_2(S, a) = \{(S', \emptyset, f, 0) \mid S' = \delta(S, a), f \text{ is } S'\text{-tight}\}$, and
 - $\delta_3 : Q_2 \times \Sigma \rightarrow 2^{Q_2}$ such that $(S', O', f', i') \in \delta_3((S, O, f, i), a)$ iff $S' = \delta(S, a)$, $f' \leq_a^S f$, $\text{rank}(f) = \text{rank}(f')$, f' is S' -tight, and
 - * $i' = (i + 2) \bmod (\text{rank}(f') + 1)$ and $O' = f'^{-1}(i')$ if $O = \emptyset$ or
 - * $i' = i$ and $O' = \delta(O, a) \cap f'^{-1}(i)$ if $O \neq \emptyset$, and
- $F' = \{\emptyset\} \cup ((2^Q \times \{\emptyset\} \times \mathcal{T} \times \omega) \cap Q_2)$.

Intuitively, COMP_S is an extension of the classical subset construction for determinization of finite automata. In particular, Q_1, δ_1 , and I_1 constitute the deterministic finite automaton obtained from \mathcal{A} using the subset construction. The automaton can, however, nondeterministically guess a point at which it will make a transition to a *macrostate* (S, O, f, i) in the Q_2 part; this guess corresponds to a level in the run DAG of the accepted word from which the ranks of all levels form an S -tight ranking, where the S component of the macrostate is again a subset from the subset construction. In the Q_2 part, \mathcal{B}_S makes sure that in order for a word to be accepted by \mathcal{B}_S , all runs of \mathcal{A} over the word need to touch an accepting state only finitely many times. This is ensured by the f component, which, roughly speaking, maps states to ranks of corresponding vertices in the run DAG over the given word. The O component is used for a standard cut-point construction, and is used to make sure that all runs that have reached an accepting state in \mathcal{A} will eventually leave it (this can happen for different runs at a different point). The S, O , and f components were already present in [24]. The i component was introduced by Schewe to improve the complexity of the construction; it is used to cycle over phases, where in each phase we focus on cut-points of a different rank. See [33] for a more elaborate exposition.

Proposition 1 (Corollary 3.3 in [33]). $\mathcal{L}(\mathcal{B}_S) = \overline{\mathcal{L}(\mathcal{A})}$.

4 Purging Macrostates with Incompatible Rankings

Our first optimisation is based on removing from \mathcal{B}_S macrostates $(S, O, f, i) \in Q_2$ whose level ranking f assigns some states of S an unnecessarily high rank. Intuitively, when S contains a state p and a state q such that p is (directly) simulated by q , i.e. $p \preceq_{di} q$, then $f(p)$ needs to be at most $f(q)$. This is because in any word α and its run DAG \mathcal{G}_α in \mathcal{A} , if p and q are at the same level i of \mathcal{G}_α , then the ranks of their vertices v_p and v_q at the given level are either both ω (when $\alpha \in \mathcal{L}(\mathcal{A})$), or such that $\text{rank}_\alpha(v_p) \leq \text{rank}_\alpha(v_q)$ otherwise. This is because, intuitively, the DAG rooted in v_p in \mathcal{G}_α is isomorphic to a subgraph of the DAG rooted in v_q .

Formally, consider the following predicate on macrostates of \mathcal{B}_S :

$$\mathcal{P}_{di}(S, O, f, i) \quad \text{iff} \quad \exists p, q \in S : p \preceq_{di} q \wedge f(p) > f(q). \quad (1)$$

We modify COMP_S to purge macrostates that satisfy \mathcal{P}_{di} . That is, we create a new procedure PURGE_{di} obtained from COMP_S by modifying the definition of \mathcal{B}_S such that all occurrences of Q_2 are substituted by Q_2^{di} and

$$Q_2^{di} = Q_2 \setminus \{(S, O, f, i) \in Q_2 \mid \mathcal{P}_{di}(S, O, f, i)\}. \quad (2)$$

We denote the BA obtained from PURGE_{di} as \mathcal{B}_S^{di} . The following lemma, proved in Section 4.1 states the correctness of this construction.

Lemma 2. $\mathcal{L}(\mathcal{B}_S^{di}) = \mathcal{L}(\mathcal{B}_S)$

The following natural question arises: Is it possible to extend the purging technique from direct simulation to other notions of simulation? For *fair* simulation, this cannot be done. The reason is that, for a pair of states p and q s.t. $p \preceq_f q$, it can happen that for a word $\beta \in \Sigma^\omega$, there can be a trace from p over β that finitely many times touches an accepting state (i.e., a vertex of p in the corresponding run DAG can have any rank between 0 and $2n$), while all traces from q over β can completely avoid touching any accepting state. From the point of view of fair simulation, these are both unfair traces, and, therefore, disregarded.

On the other hand, *delayed* simulation—which is often much richer than direct simulation—can be used, with a small change. Intuitively, the delayed simulation can be used because $p \preceq_{de} q$ guarantees that on every level of trees in \mathcal{G}_α rooted in v_p and in v_q respectively, the rank of the vertex v_p is at most by one larger than the rank of vertex v_q (or by any number smaller). Formally, let \mathcal{P}_{de} be the following predicate on macrostates of \mathcal{B}_S :

$$\mathcal{P}_{de}(S, O, f, i) \quad \text{iff} \quad \exists p, q \in S : p \preceq_{de} q \wedge f(p) > \llbracket f(q) \rrbracket, \quad (3)$$

where $\llbracket x \rrbracket$ for $x \in \omega$ denotes the smallest even number greater or equal to x and $\llbracket \omega \rrbracket = \omega$. Similarly as above, we create a new procedure, called PURGE_{de} , which is obtained from COMP_S by modifying the definition of \mathcal{B}_S such that all occurrences of Q_2 are substituted by Q_2^{de} and

$$Q_2^{de} = Q_2 \setminus \{(S, O, f, i) \in Q_2 \mid \mathcal{P}_{de}(S, O, f, i)\}. \quad (4)$$

We denote the BA obtained from PURGE_{de} as \mathcal{B}_S^{de} .

Lemma 3. $\mathcal{L}(\mathcal{B}_S^{de}) = \mathcal{L}(\mathcal{B}_S)$

The use of $\llbracket f(q) \rrbracket$ in \mathcal{P}_{de} results in the fact that the two purging techniques are incomparable. For instance, consider a macrostate $(\{p, q\}, \emptyset, \{p \mapsto 2, q \mapsto 1\}, 0)$ such that $p \preceq_{di} q$ and $p \preceq_{de} q$. Then the macrostate will be purged in PURGE_{di} , but not in PURGE_{de} .

The two techniques can, however, be easily combined into a third procedure PURGE_{di+de} , when Q_2 is substituted in COMP_S with Q_2^{di+de} defined as

$$Q_2^{di+de} = Q_2 \setminus \{(S, O, f, i) \in Q_2 \mid \mathcal{P}_{di}(S, O, f, i) \vee \mathcal{P}_{de}(S, O, f, i)\}. \quad (5)$$

We denote the resulting BA as \mathcal{B}_S^{di+de} .

Lemma 4. $\mathcal{L}(\mathcal{B}_S^{di+de}) = \mathcal{L}(\mathcal{B}_S)$

4.1 Proofs of Lemmas 2, 3, and 4

We first give a lemma that an x -strategy σ_x preserves an x -simulation \preceq_x .

Lemma 5. *Let \preceq_x be an x -simulation (for $x \in \{di, de, f\}$). Then, the following holds: $\forall p, q \in Q : p \preceq_x q \wedge p \xrightarrow{a} p' \in \delta \Rightarrow \exists q' \in Q : q \xrightarrow{a} q' \in \delta \wedge p' \preceq_x q'$.*

Proof. Let $p, q \in Q$ such that $p \preceq_x q$ and $p \xrightarrow{a} p' \in \delta$, and let π_p be a trace starting from p with the first transition $p \xrightarrow{a} p'$. From the definition of x -simulation, there is a winning Duplicator strategy σ_x ; let $\pi_q = \sigma_x(q', \pi_p)$ and let $q \xrightarrow{a} q'$ be the first transition of π_q . Let $\pi_{p'}$ and $\pi_{r'}$ be traces obtained from π_p and π_r by removing their first transitions. It is easy to see that if $\mathcal{C}^x(\pi_p, \pi_r)$ then also $\mathcal{C}^x(\pi_{p'}, \pi_{r'})$ for any $x \in \{di, de, f\}$. It follows that σ_x is also a winning Duplicator strategy from (p', r') . \square

Next, we focus on delayed simulation and the proof of Lemma 3. In the next lemma, we show that if there is a pair of vertices on some level of the run DAG where one vertex delay-simulates the other one, there exists a relation between their rankings. This will be used to purge some useless rankings from the complemented BA.

Lemma 6. *Let $p, q \in Q$ such that $p \preceq_{de} q$ and $\mathcal{G}_\alpha = (V, E)$ be the run DAG of \mathcal{A} over α . For all $i \geq 0$, it holds that $(p, i) \in V \wedge (q, i) \in V \Rightarrow \text{rank}_\alpha(p, i) \leq \llbracket \text{rank}_\alpha(q, i) \rrbracket$.*

Proof. Consider some $(p, i) \in V$ and $(q, i) \in V$. First, suppose that $\text{rank}_\alpha(q, i) = \omega$. Since the rank can be at most ω , it will always hold that $\text{rank}_\alpha(p, i) \leq \llbracket \text{rank}_\alpha(q, i) \rrbracket$.

On the other hand, suppose that $\text{rank}_\alpha(q, i)$ is finite, i.e., $\alpha_{i:\omega}$ is not accepted by q . Then, due to Lemma 1, $0 \leq \text{rank}_\alpha(q, i) \leq 2n$. Because $p \preceq_{de} q$, it holds that $\alpha_{i:\omega}$ is also not accepted by p , and therefore also $0 \leq \text{rank}_\alpha(p, i) \leq 2n$. We now need to show that $0 \leq \text{rank}_\alpha(p, i) \leq \llbracket \text{rank}_\alpha(q, i) \rrbracket \leq 2n$.

Let $\{\mathcal{G}_\alpha^k\}_{k=0}^{2n+1}$ be the sequence of run DAGs obtained from \mathcal{G}_α in the ranking procedure from Section 2.2. In the following text we use the abbreviation $v \in \mathcal{G}_\alpha^m \setminus \mathcal{G}_\alpha^n$ for $v \in \mathcal{G}_\alpha^m \wedge v \notin \mathcal{G}_\alpha^n$. Since the rank of a node (r, j) is given as the number l s.t. $(r, j) \in \mathcal{G}_\alpha^l \setminus \mathcal{G}_\alpha^{l+1}$, we will finish the proof of this lemma by proving the following claim:

Claim. Let k and l be s.t. $(p, i) \in \mathcal{G}_\alpha^k \setminus \mathcal{G}_\alpha^{k+1}$ and $(q, i) \in \mathcal{G}_\alpha^l \setminus \mathcal{G}_\alpha^{l+1}$. Then $k \leq \llbracket l \rrbracket$.

Proof: We prove the claim by induction on l .

- Base case: ($l = 0$) Since we assume \mathcal{A} is complete, no vertex in \mathcal{G}_α^0 is finite. ($l = 1$) We prove that if (q, i) is endangered in \mathcal{G}_α^1 , then (p, i) is endangered in \mathcal{G}_α^1 as well (so both would be removed in \mathcal{G}_α^2). For the sake of contradiction, assume that (q, i) is endangered in \mathcal{G}_α^1 and (p, i) is not. Therefore, since \mathcal{G}_α^1 contains no finite vertices, there is an infinite path π from (p, i) s.t. π contains at least one accepting state. In the following, we abuse notation and, given a strategy σ_{de} and a state $s \in Q$, use $\sigma_{de}((s, i), \pi)$ to denote the path $(s_0, i)(s_1, i+1)(s_2, i+2) \dots$ such that $s_0 = s$ and $\forall j \geq 0$, it holds that $s_{j+1} = \sigma_{de}(s_j, r_{i+j} \xrightarrow{\alpha_{i+j}} r_{i+j+1})$ where $\pi_x = (r_x, x)$ for every $x \geq 0$. Since $p \preceq_{de} q$, there is a corresponding infinite path $\pi' = \sigma_{de}((q, i), \pi)$ that also contains at least one accepting state. Therefore, (q, i) is not endangered, a contradiction to the assumption, so we conclude that $l = 1 \Rightarrow k = 1$.
- Inductive step: We assume the claim holds for all $l < 2j$ and prove the inductive step for even and odd steps independently.

($l = 2j$) We prove that if (q, i) is finite in \mathcal{G}_α^l (and therefore would be removed in \mathcal{G}_α^{l+1}), then either $(p, i) \notin \mathcal{G}_\alpha^l$, or (p, i) is also finite in \mathcal{G}_α^l . For the sake of contradiction, we assume that (q, i) is finite in \mathcal{G}_α^l and that (p, i) is in \mathcal{G}_α^l , but is not finite there (and, therefore, $k > l$). Since (p, i) is not finite in \mathcal{G}_α^l , there is an infinite path π from (p, i) in \mathcal{G}_α^l . Because $p \preceq_{de} q$, it follows that there is an infinite path $\pi' = \sigma_{de}((q, i), \pi)$ in \mathcal{G}_α^0 (π' is not in \mathcal{G}_α^l because (q, i) is finite there). Using Lemma 5 (possibly multiple times) and the fact that (q, i) is finite, we can find vertices (p', x) in π and (q', x) in π' s.t. $p' \preceq_{de} q'$ and (q', x) is not in \mathcal{G}_α^l , therefore, $(q', x) \in \mathcal{G}_\alpha^e \setminus \mathcal{G}_\alpha^{e+1}$ for some $e < l$. Because $(p', x) \in \mathcal{G}_\alpha^l$ and it is not finite (π is infinite), it follows that $(p', x) \in \mathcal{G}_\alpha^f \setminus \mathcal{G}_\alpha^{f+1}$ for some $f > l$, and since $e < l < f$, we have that $f \not\leq e + 1$, implying $f \not\leq \llbracket e \rrbracket$, which is in contradiction to the induction hypothesis.

($l = 2j + 1$) We prove that if (q, i) is endangered in \mathcal{G}_α^l (and therefore would be removed in \mathcal{G}_α^{l+1}), then either $(p, i) \notin \mathcal{G}_\alpha^l$, or (p, i) is removed at the latest in \mathcal{G}_α^{l+1} . For the sake of contradiction, assume that (q, i) is endangered in \mathcal{G}_α^l while (p, i) is removed later than in \mathcal{G}_α^{l+1} . Therefore, since \mathcal{G}_α^l contains no finite vertices (they were removed in the $(l-1)$ -th step), there is an infinite path π from (p, i) s.t. π contains at least one accepting state. Because $p \preceq_{de} q$, there is a corresponding path $\pi' = \sigma_{de}((q, i), \pi)$ from (q, i) in \mathcal{G}_α^0 that also contains at least one accepting state and moreover $\pi' \notin \mathcal{G}_\alpha^l$. Since π' has an infinite number of states (and at least one accepting), not all states from π' were removed in \mathcal{G}_α^{l-1} , i.e., there is at least one node with rank less or equal to $l - 2$. Using Lemma 5 (also possibly multiple times) we can hence find states (p', x) in π and (q', x) in π' s.t. $p' \preceq_{de} q'$ and (q', x) is not in \mathcal{G}_α^l and has a rank less or equal to $l - 2$, therefore, $(q', x) \in \mathcal{G}_\alpha^e \setminus \mathcal{G}_\alpha^{e+1}$ for some $e < l - 1$. Because $(p', x) \in \mathcal{G}_\alpha^l$, it follows that $(p', x) \in \mathcal{G}_\alpha^f \setminus \mathcal{G}_\alpha^{f+1}$ for some $f \geq l$, and, therefore, $f \not\leq e + 1$, which is in contradiction to the induction hypothesis. \blacksquare

This concludes the proof. \square

Lemma 7. *Let $p, q \in Q$ such that $p \preceq_{di} q$ and $\mathcal{G}_\alpha = (V, E)$ be the run DAG of \mathcal{A} over α . For all $i \geq 0$, it holds that $(p, i) \in V \wedge (q, i) \in V \Rightarrow \text{rank}_\alpha(p, i) \leq \text{rank}_\alpha(q, i)$.*

Proof. Can be obtained as a simplified version of the proof of Lemma 6. \square

We are now ready to prove Lemma 3.

Lemma 3. $\mathcal{L}(\mathcal{B}_S^{de}) = \mathcal{L}(\mathcal{B}_S)$

Proof. (\subseteq) Follows directly from the fact that \mathcal{B}_S^{de} is obtained by removing states from \mathcal{B}_S .

(\supseteq) Let $\alpha \in \mathcal{L}(\mathcal{B}_S)$. As shown in the proof of Lemma 3.2 in [33], there are two cases. The first case is when all vertices of \mathcal{G}_α are finite, which we do not need to consider, since we assume complete automata.

The other case is when \mathcal{G}_α contains an infinite vertex. In this case, \mathcal{B}_S contains an accepting run

$$\rho = S_0 S_1 \dots S_p (S_{p+1}, O_{p+1}, f_{p+1}, i_{p+1}) (S_{p+2}, O_{p+2}, f_{p+2}, i_{p+2}) \dots$$

with

- $S_0 = I, O_{p+1} = \emptyset$, and $i_{p+1} = 0$,
 - $S_{j+1} = \delta(S_j, \alpha_j)$ for all $j \in \omega$,
- and, for all $j > p$,
- $O_{j+1} = f_{j+1}^{-1}(i_{j+1})$ if $O_j = \emptyset$ or $O_{j+1} = \delta(O_j, \alpha_j) \cap f_{j+1}^{-1}(i_{j+1})$ if $O_j \neq \emptyset$, respectively,
 - f_j is the S_j -tight level ranking that maps each $q \in S_j$ to the rank of $(q, j) \in \mathcal{G}_\alpha$,
 - $i_{j+1} = i_j$ if $O_j \neq \emptyset$ or $i_{j+1} = (i_j + 2) \bmod (\text{rank}(f) + 1)$ if $O_j = \emptyset$, respectively.

The ranks assigned by f_j to states of S_j match the ranks of the corresponding vertices in \mathcal{G}_α .

⊗ Using Lemma 6, we conclude that ρ contains no macrostate (S, O, f, j) where $f(p) > \llbracket f(q) \rrbracket$ and $p \preceq_{de} q$ for $p, q \in S$. Therefore, ρ is also an accepting run in \mathcal{B}_S^{de} . (We use ⊗ to refer to this paragraph later.) \square

Lemma 2. $\mathcal{L}(\mathcal{B}_S^{di}) = \mathcal{L}(\mathcal{B}_S)$

Proof. The same as for Lemma 3 with ⊗ substituted by the following:

⊗ Using Lemma 7, we conclude that ρ contains no macrostate (S, O, f, j) where $f(p) > f(q)$ and $p \preceq_{di} q$ for $p, q \in S$. So ρ is also an accepting run in \mathcal{B}_S^{di} . \square

Lemma 4. $\mathcal{L}(\mathcal{B}_S^{di+de}) = \mathcal{L}(\mathcal{B}_S)$

Proof. The same as for Lemma 3 with ⊗ substituted by the following:

⊗ Using Lemmas 7 and 6, we conclude that ρ contains no macrostate (S, O, f, j) where either $f(p) > f(q)$ and $p \preceq_{di} q$, or $f(p) > \llbracket f(q) \rrbracket$ and $p \preceq_{de} q$ for $p, q \in S$. Therefore, ρ is also an accepting run in \mathcal{B}_S^{di+de} . \square

5 Saturation of Macrostates

Our second optimisation is inspired by an optimisation of determinisation of classical finite automata from [17, Section 5]. Their optimisation is based on saturating every constructed macrostate in the classical subset construction with all direct-simulation-smaller states. This can reduce the total number of states of the determinized automaton because two or more macrostates can be mapped to a single saturated macrostate. (In Section 5.2, we show why an analogue of their *compression* cannot be used.)

We show that a similar technique can be applied to BAs. We do not restrain ourselves to direct simulation, though, and generalize the technique to delayed simulation. In particular, in our optimisation, we saturate the S components of macrostates (S, O, f, i) obtained in COMP_S with all \preceq_{de} -smaller states. Formally, we modify COMP_S by substituting the definition of the constructed transition function δ' with δ'_{Sat} defined as follows:

- $\delta'_{Sat} = \delta_1^{Sat} \cup \delta_2^{Sat} \cup \delta_3^{Sat}$ where
 - $\delta_1^{Sat} : Q_1 \times \Sigma \rightarrow 2^{Q_1}$ with $\delta_1^{Sat}(S, a) = \{cl[\delta(S, a)]\}$,
 - $\delta_2^{Sat} : Q_1 \times \Sigma \rightarrow 2^{Q_2}$ with $\delta_2^{Sat}(S, a) = \{(S', \emptyset, f, 0) \mid S' = cl[\delta(S, a)]\}$,
and
 - $\delta_3^{Sat} : Q_2 \times \Sigma \rightarrow 2^{Q_2}$ with $(S', O', f', i') \in \delta_3^{Sat}((S, O, f, i), a)$ iff $S' = cl[\delta(S, a)]$, $f' \leq_a^S f$, $rank(f) = rank(f')$, and
 - * $i' = (i + 2) \bmod (rank(f') + 1)$ and $O' = f'^{-1}(i')$ if $O = \emptyset$ or
 - * $i' = i$ and $O' = \delta(O, a) \cap f'^{-1}(i)$ if $O \neq \emptyset$,

where $cl[S] = \{q \in Q \mid \exists s \in S : q \preceq_{de} s\}$. We denote the obtained procedure as SATURATE and the obtained BA as \mathcal{B}_{Sat} .

Lemma 8. $\mathcal{L}(\mathcal{B}_{Sat}) = \mathcal{L}(\mathcal{B}_S)$

Obviously, as direct simulation is stronger than delayed simulation, the previous technique can also use direct simulation only (e.g., when computing the full delayed simulation is computationally too demanding). Moreover, SATURATE is also compatible with all PURGE_x algorithms for $x \in \{di, de, di + de\}$ (because they just remove macrostates with incompatible rankings from Q_2)—we call the combined versions PURGE_x+SATURATE and the complement BAs they output \mathcal{B}_{Sat}^x .

Lemma 9. $\mathcal{L}(\mathcal{B}_{Sat}^{di}) = \mathcal{L}(\mathcal{B}_{Sat}^{de}) = \mathcal{L}(\mathcal{B}_{Sat}^{di+de}) = \mathcal{L}(\mathcal{B}_S)$

5.1 Proofs of Lemmas 8 and 9

We start with a lemma, used later, that talks about languages of states related by delayed simulation when there is a path between them.

Lemma 10. For $p, q \in Q$ such that $p \preceq_{de} q$, let $L_\top = \{w \in \Sigma^* \mid p \xrightarrow{w}_F q\}$ and $L_\perp = \{w \in \Sigma^* \mid p \xrightarrow{w} q\}$. Then $L(q) \supseteq (L_\perp^* L_\top)^\omega$.

Proof. First we prove the following claim:

Claim. For every word $\alpha = w_0 w_1 w_2 \dots \in \Sigma^\omega$ where $w_i \in L_\top \cup L_\perp$, we can construct a trace $\pi = p \xrightarrow{w_0} q_0 \xrightarrow{w_1} q_1 \xrightarrow{w_2} \dots$ over α such that $p \preceq_{de} q_0$ and $q_i \preceq_{de} q_{i+1}$ for all $i \geq 0$.

Proof: We assign $q_0 := p$ and construct the rest of π by the following inductive construction.

- Base case: ($i = 0$) From the assumption it holds that $p \xrightarrow{w_1} q_0$ and $p \preceq_{de} q_0$. From Lemma 5 there is some $r \in Q$ s.t. $q_0 \xrightarrow{w_1} r$ and $q_0 \preceq_{de} r$. We assign $q_1 := r$, so $q_0 \preceq_{de} q_1$.
- Inductive step: Let $\pi' = p \xrightarrow{w_0} q_0 \xrightarrow{w_1} \dots \xrightarrow{w_i} q_i$ be a prefix of a trace such that $q_j \preceq_{de} q_{j+1}$ for every $j < i$. From the transitivity of \preceq_{de} , it follows that $p \preceq_{de} q_i$. From Lemma 5 there is some $r \in Q$ s.t. $q_i \xrightarrow{w_{i+1}} r$ and $q_i \preceq_{de} r$. We assign $q_{i+1} := r$, so $q_i \preceq_{de} q_{i+1}$. ■

Consider a word $\alpha \in (L_\perp^* L_\top)^\omega$ such that $\alpha = w_0 w_1 w_2 \dots$ for $w_i \in L_\top \cup L_\perp$. We show that $\alpha \in \mathcal{L}(q)$. According to the previous claim, we can construct a trace $\pi = p \xrightarrow{w_0} q = q_0 \xrightarrow{w_1} q_1 \xrightarrow{w_2} \dots$ over α s.t. $p \preceq_{de} q_0$ and $q_i \preceq_{de} q_{i+1}$ for all $i \geq 0$. Since $p \preceq_{de} q$, from Lemma 5 it follows that we can construct a trace $\pi' = q \xrightarrow{w_0} r_0 \xrightarrow{w_1} r_1 \xrightarrow{w_2} \dots$ s.t. $q_i \preceq_{de} r_i$ for every $i \geq 0$. Because α contains infinitely often a subword from L_\top , there is some $\ell \in \omega$ such that $q_\ell \xrightarrow{w_\ell} q_{\ell+1}$ and $r_\ell \xrightarrow{w_\ell} r_{\ell+1}$ for $w_\ell \in L_\top$. Note that it holds that $p \preceq_{de} q_\ell \preceq_{de} r_\ell$. We can again use the claim above to construct a trace $\pi^* = p \xrightarrow[F]{w_\ell} q = s_0 \xrightarrow{w_{\ell+1}} s_1 \xrightarrow{w_{\ell+2}} \dots$ over $\alpha_\ell = w_\ell w_{\ell+1} w_{\ell+2} \dots$ such that $p \preceq_{de} s_0$ and $s_i \preceq_{de} s_{i+1}$ for all $i \geq 0$. Since $p \preceq_{de} r_\ell$, we can simulate π^* from r_ℓ by a trace $\pi^{*'}$, and because $p \xrightarrow[F]{w_\ell} q$, we know that $\pi^{*'}$ will touch an accepting state in finitely many steps (this holds because w_ℓ is from L_\top , which are the words over which we can go from p to q and touch an accepting state). Consider $m \geq \ell$ such that s_m is the first state after the accepting state that is one of the $\{s_0, s_1, \dots\}$ in $\pi^{*'}$. This reasoning could be repeated for all occurrences of a subword from L_\top in π^* , therefore $\alpha \in \mathcal{L}(q)$. \square

Next, we give a lemma used for establishing correctness of saturating macrostates with \preceq_{de} -smaller states.

Lemma 11. *Let $p, q, r \in Q$ such that $r \xrightarrow{a} q \in \delta$ and $p \preceq_{de} q$. Further, let $\mathcal{A}' = (Q, \delta', I, F)$ where $\delta' = \delta \cup \{r \xrightarrow{a} p\}$. Then $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$.*

Proof. (\subseteq) Clear.

(\supseteq) Consider some $\alpha \in \mathcal{L}(\mathcal{A}')$ and an accepting trace π in \mathcal{A}' over α . There are two cases:

1. (π contains only finitely many transitions $r \xrightarrow{a} p$)

In this case, π is of the form $\pi = \pi_i \pi_\omega$ where π_i is a finite prefix $\pi_i = q_0 \xrightarrow{w_0} r \xrightarrow{a} p \xrightarrow{w_1} r \xrightarrow{a} p \xrightarrow{w_2} \dots \xrightarrow{w_n} r \xrightarrow{a} p$, for $q_0 \in I$, and π_ω is an infinite trace from p that does not contain any occurrence of the transition $r \xrightarrow{a} p$. We construct in \mathcal{A} a trace $\pi' = q_0 \xrightarrow{w_0} r \xrightarrow{a} q \xrightarrow{w_1} r_1 \xrightarrow{a} q_1 \xrightarrow{w_2} \dots \xrightarrow{w_n} r_n \xrightarrow{a} q_n \cdot \pi'_\omega$ as follows. Let σ_{de} be a strategy for \preceq_{de} . We set $r_1 := \sigma_{de}(q, p \xrightarrow{w_1} r)$, so $r \preceq_{de} r_1$. Since $r \xrightarrow{a} q \in \delta$, it follows that there is $r_1 \xrightarrow{a} q_1 \in \delta$ such that $p \preceq_{de} q_1$. For $i > 1$, we set $r_i := \sigma_{de}(q_{i-1}, p \xrightarrow{w_i} r)$. By induction, it follows that $\forall 1 \leq i \leq n : p \preceq_{de} q_i$, in particular $p \preceq_{de} q_n$. We set $\pi'_\omega := \sigma_{de}(q_n, \pi_\omega)$. Since π_ω starts in p and contains infinitely many accepting states and π'_ω starts in q_n and $p \preceq_{de} q_n$, then π'_ω also contains infinitely many accepting states. It follows that π' is accepting, so $\alpha \in \mathcal{L}(\mathcal{A})$.

2. (π contains infinitely many transitions $r \xrightarrow{a} p$)

In this case, π is of the form $\pi = q_0 \xrightarrow{w_0} r \xrightarrow{a} p \xrightarrow{w_1} r \xrightarrow{a} p \xrightarrow{w_2} \dots \xrightarrow{w_n} r \xrightarrow{a} p \xrightarrow{w_{n+1}} \dots$, for $q_0 \in I$ and $\alpha = w_0 a w_1 a w_2 \dots$. Since π is accepting, for infinitely many $i \in \omega$, we have $p \xrightarrow[F]{w_i a} p$ in \mathcal{A}' and hence also $p \xrightarrow[F]{w_i a} q$ in the original BA \mathcal{A} . Using Lemma 10 and the fact that $p \preceq_{de} q$, we have $w_1 a w_2 a \dots \in L(q)$ and hence $\alpha = w_0 a w_1 a w_2 a \dots \in \mathcal{L}(\mathcal{A})$. \square

The following lemma guarantees that adding transitions in the way of Lemma 11 does not break the computed delayed simulation and can, therefore, be performed repeatedly, without the need to recompute the simulation.

Lemma 12. Let \preceq_{de} be the delayed simulation on \mathcal{A} . Further, let $p, q, r \in Q$ be such that $r \xrightarrow{a} q \in \delta$ and $p \preceq_{de} q$, and let $\mathcal{A}' = (Q, \delta', I, F)$ where $\delta' = \delta \cup \{r \xrightarrow{a} p\}$. Then \preceq_{de} is included in the delayed simulation on \mathcal{A}' .

Proof. Let σ_{de} be a dominating strategy compatible with \preceq_{de} and σ'_{de} be a strategy defined for all $s \in Q$ such that $r \preceq_{de} s$ as $\sigma'_{de}(s, x) = \sigma_{de}(s, x)$ when $x \neq (r \xrightarrow{a} p)$ and $\sigma'_{de}(s, r \xrightarrow{a} p) = \sigma_{de}(s, r \xrightarrow{a} q)$. Note that σ'_{de} is also dominating wrt \preceq_{de} . This can be shown by the following proof by contradiction: Suppose σ'_{de} is not dominating; then there is a strategy ρ such that $\sigma'_{de}(s, r \xrightarrow{a} p)$ must be simulated by $\rho(s, r \xrightarrow{a} p) = t$. But then $\sigma_{de}(s, r \xrightarrow{a} q)$ must also (transitivity of simulation) be simulated by t , so σ_{de} is not dominating. Contradiction.

Further, let $t, u \in Q$ be such that $t \preceq_{de} u$. Let $\pi_t = t \xrightarrow{w_1} t_f \xrightarrow{w_2} r \xrightarrow{a} p. \pi'_t$ be a trace over $\alpha = w_1 w_2 a w_\omega \in \Sigma^\omega$ in \mathcal{A}' such that t_f is an accepting state and $t_f \xrightarrow{w_2} r$ does not contain any occurrence of $r \xrightarrow{a} p$. Further, let $\pi_u = u_0 \xrightarrow{w_1} u_f \xrightarrow{w_2} u_i \xrightarrow{a} u_{i+1}. \pi'_u$ be a trace corresponding to a run $u_0 u_1 u_2 \dots$ over α in \mathcal{A} , where $u_0 = u$, constructed as $\pi_u = \sigma'_{de}(u, \pi_t)$.

Claim. There is a trace $\pi_v = t \xrightarrow{w_1} v_f. \pi'_v$ over α such that π'_v contains an accepting state and π_v is \preceq_{de} -simulated by π_u at every position.

Proof: We have the following two cases:

- $(t \xrightarrow{w_1} t_f \text{ does not contain any occurrence of } r \xrightarrow{a} p)$
 Let $\pi_v = t \xrightarrow{w_1} t_f \xrightarrow{w_2} r \xrightarrow{a} q. \pi'_v$ be a trace in \mathcal{A} over α obtained from π_t by starting with its prefix up to r , taking $r \xrightarrow{a} q$, and continuing with $\pi'_v = \sigma'_{de}(q, \pi'_t)$. Since in π_v , it holds that t_f is at the same position as t_f in π_t , the first part of the claim holds. Further, π_u clearly \preceq_{de} -simulates π_v on $t \xrightarrow{w_1} t_f \xrightarrow{w_2} r$, and because σ'_{de} simulates $r \xrightarrow{a} p$ by a transition to a state u_{i+1} such that $q \preceq_{de} u_{i+1}$ and π'_v is constructed using σ'_{de} , then also the second part of the claim holds.
- $(t \xrightarrow{w_1} t_f \text{ contains at least one occurrence of } r \xrightarrow{a} p)$
 Suppose that π_t starts with $t \xrightarrow{w_{11}} r \xrightarrow{a} p \xrightarrow{w_{12}} t_f$ such that $t \xrightarrow{w_{11}} r$ does not contain any $r \xrightarrow{a} p$. Then let us start building π_v such that it starts with $t \xrightarrow{w_{11}} r \xrightarrow{a} q$. On this prefix, π_v is clearly \preceq_{de} -simulated by the corresponding prefix of π_u . We continue from q using the strategy σ'_{de} . In particular, the next time we reach $r \xrightarrow{a} p$ in π_t while we are at some state v_1 such that $r \preceq_{de} v_1$, we simulate the transition by $\sigma'_{de}(v_1, r \xrightarrow{a} p)$ and so on. We can observe that when we arrive to t_f in π_t , we also arrive to v_f in π_v such that $t_f \preceq_{de} v_f$. Therefore, π'_v contains an accepting state. Moreover, since σ'_{de} is dominating, the second part of the claim also holds. ■

From the claim above, it follows that the trace $u_f \xrightarrow{w_2} u_i \xrightarrow{a} u_{i+1}. \pi'_u$ contains an accepting state, so $\mathcal{C}^{de}(\pi_t, \pi_u)$. □

Finally, we are ready to prove Lemma 8.

Lemma 8. $\mathcal{L}(\mathcal{B}_{Sat}) = \mathcal{L}(\mathcal{B}_S)$

- Proof.* (\subseteq) Let $\alpha \in \mathcal{L}(\mathcal{B}_{Sat})$ and ρ be an arbitrary accepting run over α in \mathcal{B}_{Sat} such that $\rho = S_0 S_1 \dots S_{n-1}(S_n, O_n, f_n, i_n)(S_{n+1}, O_{n+1}, f_{n+1}, i_{n+1}) \dots$. For the sake of contradiction, assume that $\alpha \in \mathcal{L}(\mathcal{A})$, therefore, there is a run ρ' on α in \mathcal{A} having infinitely many accepting states. From the fact that tight level rankings form a non-increasing sequence, we have that $f_n(\rho'(n)) \geq f_{n+1}(\rho'(n+1)) \geq \dots$. This sequence eventually stabilizes and from the property of level rankings and the fact that ρ' is accepting, it stabilizes in some ℓ such that $f_\ell(\rho'(\ell))$ is even. This, however, means that the O component of macrostates in ρ cannot be emptied infinitely often, and, therefore, ρ is not accepting, which is a contradiction. Hence $\alpha \notin \mathcal{L}(\mathcal{A})$, so (from Proposition 1) $\alpha \in \mathcal{L}(\mathcal{B}_S)$.
- (\supseteq) Consider some $\alpha \in \mathcal{L}(\mathcal{B}_S)$. Let \mathcal{A}' be a BA obtained from \mathcal{A} by adding transitions according to Lemma 12. Then from Lemma 11, we have that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$. Therefore, $\alpha \in \mathcal{L}(\mathcal{B}'_S)$ where \mathcal{B}'_S is the BA obtained from \mathcal{A}' using COMPs. It is easy to see that we can construct a run in \mathcal{B}_{Sat} that mimics the levels of run DAG of α in \mathcal{A}' (i.e., we are able to empty the O component infinitely often). Hence $\alpha \in \mathcal{L}(\mathcal{B}_{Sat})$. \square


Lemma 9. $\mathcal{L}(\mathcal{B}_{Sat}^{di}) = \mathcal{L}(\mathcal{B}_{Sat}^{de}) = \mathcal{L}(\mathcal{B}_{Sat}^{di+de}) = \mathcal{L}(\mathcal{B}_S)$

- Proof.* (\subseteq) This part is the same as in the proof of Lemma 8.
- (\supseteq) Consider some $\alpha \in \mathcal{L}(\mathcal{B}_S)$. Let \mathcal{A}' be a BA obtained from \mathcal{A} by adding transitions according to Lemma 12. Then from Lemma 11, we have that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\mathcal{A}')$. Therefore, $\alpha \in \mathcal{L}(\mathcal{B}'_S)$ where \mathcal{B}'_S is the BA obtained from \mathcal{A}' using COMPs. It is easy to see that we can construct a run in \mathcal{B}_{Sat} that mimics the levels of run DAG of α in \mathcal{A}' (i.e., we are able to empty the O component infinitely often). Using Lemmas 7 and 6, we can conclude that the run contains no macrostate of the form (S, O, f, j) , where $f(p) > f(q)$ and $p \preceq_{di} q$, or $f(p) > \llbracket f(q) \rrbracket$ and $p \preceq_{de} q$ for $p, q \in S$. Therefore, ρ is also an accepting run in $\mathcal{B}_{Sat}^{di+de}$. Hence $\alpha \in \mathcal{L}(\mathcal{B}_{Sat}^{di+de})$. \square

5.2 Remarks on Compression of Macrostates

An analogy to saturation of macrostates is their compression [17, Section 6], based on removing simulation-smaller states from a macrostate. This is, however, not possible even for direct simulation, as we can see in the following example.

Example 1. Consider the BA over $\Sigma = \{a\}$ given below.



For this BA we have $q \preceq_{di} r$ and $r \preceq_{di} q$. If we compress the macrostates obtained in COMP_S , there is the following trace in the output automaton:

$$\begin{aligned} \{p\} &\xrightarrow{a}(\{p, q\}, \emptyset, \{p \mapsto 3, q \mapsto 2, r \mapsto 1\}, 0) \xrightarrow{a}(\{p, r\}, \{r\}, \{p \mapsto 3, q \mapsto 1, r \mapsto 2\}, 2) \\ &\xrightarrow{a}(\{p, q\}, \emptyset, \{p \mapsto 3, q \mapsto 2, r \mapsto 1\}, 2) \xrightarrow{a}(\{p, r\}, \{r\}, \{p \mapsto 3, q \mapsto 1, r \mapsto 2\}, 0) \\ &\xrightarrow{a}(\{p, q\}, \emptyset, \{p \mapsto 3, q \mapsto 2, r \mapsto 1\}, 0) \xrightarrow{a} \dots \end{aligned}$$

This trace contains infinitely many final states (we flush the O -set infinitely often), hence we are able to accept the word a^ω , which is, however, in the language of the input BA. \square

6 Use after Simulation Quotienting

In this short section, we establish that our optimizations introduced in Sections 4 and 5 can be applied with no additional cost in the setting when BA complementation is preceded with simulation-based reduction of the input BA (which is usually helpful), i.e., when the simulation is already computed beforehand for another purpose. In particular, we show that simulation-based reduction preserves the simulation (when naturally extended to the quotient automaton). First, let us formally define the operation of quotienting.

Given an x -simulation \preceq_x for $x \in \{di, de\}$, we use \approx_x to denote the x -similarity relation (i.e., the symmetric fragment) $\approx_x = \preceq_x \cap \preceq_x^{-1}$. Note that since \preceq_x is a preorder, it holds that \approx_x is an equivalence. We use $[q]_x$ to denote the equivalence class of q wrt \approx_x . The *quotient* of a BA $\mathcal{A} = (Q, \delta, I, F)$ wrt \approx_x is the automaton

$$\mathcal{A}/\approx_x = (Q/\approx_x, \delta_{\approx_x}, I_{\approx_x}, F_{\approx_x}) \quad (6)$$

with the transition function $\delta_{\approx_x}([q]_x, a) = \{[r]_x \mid r \in \delta([q]_x, a)\}$ and the set of initial and accepting states $I_{\approx_x} = \{[q]_x \in Q/\approx_x \mid q \in I\}$ and $F_{\approx_x} = \{[q]_x \in Q/\approx_x \mid q \in F\}$ respectively.

Proposition 2 ([7], [12]). *If $x \in \{di, de\}$, then $\mathcal{L}(\mathcal{A}/\approx_x) = \mathcal{L}(\mathcal{A})$.*

Remark 1 ([12]). $\mathcal{L}(\mathcal{A}/\approx_f) \neq \mathcal{L}(\mathcal{A})$

Finally, the following lemma shows that quotienting preserves direct and delayed simulations, therefore, when complementing \mathcal{A} , it is possible to first quotient \mathcal{A} wrt a direct/delayed simulation and then use the same simulation (lifted to the states of the quotient automaton) to optimize the complementation.

Lemma 13. *Let \preceq_x be the x -simulation on \mathcal{A} for $x \in \{di, de\}$. Then the relation \preceq_x^\approx defined as $[q]_x \preceq_x^\approx [r]_x$ iff $q \preceq_x r$ is the x -simulation on \mathcal{A}/\approx_x .*

Proof. First, we show that \preceq_x^\approx is well defined, i.e., if $q \preceq_x r$, then for all $q' \in [q]_x$ and $r' \in [r]_x$, it holds that $q' \preceq_x r'$. Indeed, this holds because $q' \approx_x q$ and $r \approx_x r'$, and therefore $q' \preceq_x q \preceq_x r \preceq_x r'$; the transitivity of simulation yields $q' \preceq_x r'$.

Next, let σ_x be a strategy that gives \preceq_x . Consider a trace defined as $[\pi_q]_x = [q_0]_x \xrightarrow{\alpha_0} [q_1]_x \xrightarrow{\alpha_1} \dots$ over a word $\alpha \in \Sigma^\omega$ in \mathcal{A}/\approx_x . Then,

1. for $x = di$ there is a trace $\pi_q = q'_0 \xrightarrow{\alpha_0} q'_1 \xrightarrow{\alpha_1} \dots$ in \mathcal{A} s.t. $q'_0 \in [q_0]_x$ and $q_i \preceq_x q'_i$ for $i \geq 0$. Therefore, if $[q_i]_x$ is accepting then so is q'_i ;
2. for $x = de$ there is a trace $\pi_q = q'_0 \xrightarrow{\alpha_0} q'_1 \xrightarrow{\alpha_1} \dots$ in \mathcal{A} s.t. $q'_0 \in [q_0]_x$, $q_i \preceq_x q'_i$ for $i \geq 0$ and, moreover, if $[q_i]_x$ is accepting then there is q'_k for $k \geq i$ s.t. $q'_k \in F$.

Further, let $[q_0]_x \preceq_x^\approx [r_0]_x$. Then there is a trace $\pi_r = \sigma_x(r, \pi_q) = (r = r_0) \xrightarrow{\alpha_0} r_1 \xrightarrow{\alpha_1} \dots$ simulating π_q in \mathcal{A} from r . Further, consider its projection $[\pi_r]_x = [r_0]_x \xrightarrow{\alpha_0} [r_1]_x \xrightarrow{\alpha_1} \dots$ into \mathcal{A}/\approx_x . For all $i \geq 0$, we have that $q_i \preceq_x r_i$, and therefore also $[q_i]_x \preceq_x^\approx [r_i]_x$. Since $\mathcal{C}^x(\pi_q, \pi_r)$, then also $\mathcal{C}^x([\pi_q]_x, [\pi_r]_x)$.

Finally, we show that \preceq_x^\approx is maximal. For the sake of contradiction, suppose that $[r]_x$ is x -simulating $[q]_x$ for some $q, r \in Q$ s.t. $q \not\preceq_x r$. Consider a word $\alpha \in \Sigma^\omega$ and a trace $\pi_q = (q = q_0) \xrightarrow{\alpha_0} q_1 \xrightarrow{\alpha_1} \dots$ over α in \mathcal{A} . Then there is a trace $[\pi_q]_x = [q = q_0]_x \xrightarrow{\alpha_0} [q_1]_x \xrightarrow{\alpha_1} \dots$ over α in \mathcal{A}/\approx_x . According to the assumption, there is also a trace $[\pi_r]_x = [r = r_0]_x \xrightarrow{\alpha_0} [r_1]_x \xrightarrow{\alpha_1} \dots$ such that $[\pi_r]_x$ is x -simulating $[\pi_q]_x$. But then there will also exist a trace $\pi_r = (r = r_0) \xrightarrow{\alpha_0} r'_1 \xrightarrow{\alpha_1} r'_1 \xrightarrow{\alpha_2} \dots$ such that $r_i \preceq_x r'_i$ for all $i \in \omega$ and $\mathcal{C}^x(\pi_q, \pi_r)$ (see the previous part of the proof). Therefore, since \preceq_x is maximal, we have that $q \preceq_x r$, which is in contradiction with the assumption. \square

7 Experimental Evaluation

We implemented our optimisations in a prototype tool⁴ written in Haskell and performed preliminary experimental evaluation on a set of 124 random BAs with a non-trivial language over a two-symbol alphabet generated using Tabakov and Vardi's model [37]. The parameters of input automata were set to the following bounds: number of states: 6–7, transition density: 1.2–1.3, and acceptance density: 0.35–0.5. Before complementing, the BAs were quotiented wrt the direct simulation for experiments with PURGE_{di} and the delayed simulation for experiments with PURGE_{de} and PURGE_{di+de} . The timeout was set to 300 s.

We present the results for our strongest optimizations for *outputs* of the size up to 500 states in Fig. 1. As can be seen in Fig. 1a, purging alone often significantly reduces the size of the output. The situation with saturation is, on the other hand, more complicated. In Fig. 1b, we can see that in some cases, the saturation produces even smaller BAs than only purging, on the other hand, in some cases, larger BAs are produced. This is expected, because saturating the S component of macrostates also means that more level rankings (the f component) need to be considered.

For outputs of a larger size (we had 11 of them), the results follow a similar trend, but the probability that saturation will increase the size of the result decreases. For some concrete results, for one BA, the size of the output BA decreased from 4065 (COMP_S) to 985 (PURGE_{di+de}) to 929 (PURGE_{di+de})

⁴ <https://github.com/vhavlena/ba-complement>




Fig. 1: Comparison of the number of states of complement BAs generated by COMP_S and our optimizations (lower is better)

+SATURATE), which yields a reduction to 24 %, resp. 22 %! Further, we observed that all PURGE_x methods usually give similar results, with the difference of only a few states (when PURGE_{di} and PURGE_{de} differ, PURGE_{di} usually wins over PURGE_{de}).

8 Related Work

BA complementation has a long research track. Known approaches can be roughly classified into Ramsey-based [34], determinization-based [32,29], rank-based [33], slice-based [23,39], learning-based [25], and the recently proposed subset-tuple construction [4]. Those approaches build on top of different concepts of capturing words accepted by a complement automaton. Some concepts can be translated into others, such as the slice-based approach, which can be translated to the rank-based approach [40]. Such a translation can help us get a deeper understanding of the BA complementation problem and the relationship between optimization techniques for different complementation algorithms.

Because of the high computational complexity of complementing a BA, and, consequently, also checking BA inclusion and universality (which use complementation as their component), there has been some effort to develop heuristics that help to reduce the number of explored states in practical cases. The most prominent ones are heuristics that leverage various notions of simulation relations, which often provide a good compromise between the overhead they impose and the achieved state space reduction. Direct [7,36], delayed [12], fair [12], their variants for alternating Büchi automata [16], and multi-pebble simulations [13] are the best-studied relations of this kind. Some of the relations can be used quotienting, but also for pruning transitions entering simulation-smaller states (which may cause some parts of the BA to become inaccessible). A series of results in this direction was recently developed by Clemente and Mayr [10,26,27].

Not only can the relations be used for reducing the size of the input BA, they can also be used for under-approximating inclusion of languages of states. For instance, during a BA inclusion test $\mathcal{L}(\mathcal{A}_S) \stackrel{?}{\subseteq} \mathcal{L}(\mathcal{A}_B)$, if every initial state of \mathcal{A}_S is simulated by an initial state of \mathcal{A}_B , the inclusion holds and no complementation needs to be performed. But simulations can also be used to reduce the explored state space within, e.g., the inclusion check itself, for instance in the context of Ramsey-based algorithms [1,2]. Ramsey-based complementation algorithms [34] in the worst case produce $2^{\mathcal{O}(n^2)}$ states, which is a significant gap from the lower bound of Michel [28] and Yan [41]. The Ramsey-based construction was, however, later improved by Breuers et al. [5] to match the upper bound $2^{\mathcal{O}(n \log n)}$. The way simulations are applied in the Ramsey-based approach is fundamentally different from the current work, which is based on rank-based construction. Taking universality checking as an example, the algorithm checks if the language of the complement automaton is empty. They run the complementation algorithm and the emptiness check together, on the fly, and during the construction check if a macrostate with a larger language has been produced before; if yes, then they can stop the search from the language-smaller macrostate. Note that, in contrast to our approach, their algorithm does not produce the complement automaton.

9 Conclusion and Future Work

We developed two novel optimizations of the rank-based complementation algorithm for Büchi automata that are based on leveraging direct and delayed simulation relations to reduce the number of states of the complemented automaton. The optimizations are directly usable in rank-based BA inclusion and universality checking. We conjecture that the decision problem of checking BA language inclusion might also bring another opportunities for exploiting simulation, such as in a similar manner as in [3]. Another, orthogonal, directions of future work are (i) applying simulation in other than the rank-based approach (in addition to the particular use within [1,2]), e.g., complementation based on Safra’s construction [32], which, according to our experience, often produces smaller complements than the rank-based procedure, (ii) applying our ideas within determinization constructions for BAs, and (iii) generalizing our techniques for richer simulations, such as the multi-pebble simulation [13] or various look-ahead simulations [26,27]. Since the richer simulations are usually harder to compute, it would be interesting to find the sweet spot between the overhead of simulation computation and the achieved state space reduction.

Acknowledgement We thank the anonymous reviewers for their helpful comments on how to improve the exposition in this paper. This work was supported by the Ministry of Science and Technology of Taiwan project 106-2221-E-001-009-MY3 the Czech Science Foundation project 19-24397S, the FIT BUT internal project FIT-S-17-4014, and The Ministry of Education, Youth and Sports from the National Programme of Sustainability (NPU II) project IT4Innovations excellence in science—LQ1602.

References

1. Abdulla, P.A., Chen, Y., Clemente, L., Holík, L., Hong, C.D., Mayr, R., Vojnar, T.: Simulation Subsumption in Ramsey-based Büchi Automata Universality and Inclusion Testing. In: Proc. of CAV’10. pp. 132–147. Springer (2010)
2. Abdulla, P.A., Chen, Y., Clemente, L., Holík, L., Hong, C., Mayr, R., Vojnar, T.: Advanced Ramsey-based Büchi Automata Inclusion Testing. In: Proc. of CONCUR’11. pp. 187–202. Springer (2011)
3. Abdulla, P.A., Chen, Y., Holík, L., Mayr, R., Vojnar, T.: When Simulation Meets Antichains. In: Proc. of TACAS’10. pp. 158–174. Springer (2010)
4. Allred, J.D., Ultes-Nitsche, U.: A Simple and Optimal Complementation Algorithm for Büchi Automata. In: Proc. of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science. pp. 46–55. ACM (2018)
5. Breuers, S., Löding, C., Olschewski, J.: Improved Ramsey-Based Büchi Complementation. In: Proc. of FOSSACS’12. pp. 150–164. Springer (2012)
6. Büchi, J.R.: On a Decision Method in Restricted Second Order Arithmetic. In: Proc. of International Congress on Logic, Method, and Philosophy of Science 1960. Stanford Univ. Press, Stanford (1962)
7. Bustan, D., Grumberg, O.: Simulation-based Minimization. ACM Transactions on Computational Logic **4**(2), 181–206 (2003)
8. Cécé, G.: Foundation for a Series of Efficient Simulation Algorithms. In: Proc. of LICS’17. pp. 1–12 (2017)
9. Chen, Y., Heizmann, M., Lengál, O., Li, Y., Tsai, M., Turrini, A., Zhang, L.: Advanced Automata-based Algorithms for Program Termination Checking. In: Proc. of PLDI’18. pp. 135–150. ACM (2018)
10. Clemente, L.: Büchi Automata Can Have Smaller Quotients. In: Proc. of ICALP’11. pp. 258–270. Springer (2011)
11. Dill, D.L., Hu, A.J., Wong-Toi, H.: Checking for Language Inclusion Using Simulation Preorders. In: Proc. of CAV’92. pp. 255–265. Springer (1992)
12. Etessami, K., Wilke, T., Schuller, R.: Fair Simulation Relations, Parity Games, and State Space Reduction for Büchi Automata. SIAM Journal on Computing **34**(5), 1159–1175 (2005)
13. Etessami, K.: A hierarchy of polynomial-time computable simulations for automata. In: Proc. of CONCUR’02. pp. 131–144. Springer (2002)
14. Fogarty, S., Vardi, M.Y.: Büchi complementation and size-change termination. In: Proc. of TACAS’09. pp. 16–30. Springer (2009)
15. Friedgut, E., Kupferman, O., Vardi, M.: Büchi Complementation Made Tighter. International Journal of Foundations of Computer Science **17**, 851–868 (2006)
16. Fritz, C., Wilke, T.: Simulation relations for alternating büchi automata. Theoretical Computer Science **338**(1), 275–314 (2005)
17. Glabbeek, R., Ploeger, B.: Five determinisation algorithms. In: Proc. of CIAA’08. pp. 161–170. Springer (2008)
18. Gurumurthy, S., Bloem, R., Somenzi, F.: Fair simulation minimization. In: Proc. of CAV’02. pp. 610–623. Springer (2002)
19. Heizmann, M., Hoenicke, J., Podelski, A.: Termination analysis by learning terminating programs. In: Proc. of CAV’14. pp. 797–813. Springer (2014)
20. Henzinger, M.R., Henzinger, T.A., Kopke, P.W.: Computing simulations on finite and infinite graphs. In: Proc. of FOCS’95. pp. 453–462. IEEE Computer Society (1995)

21. Henzinger, T.A., Kupferman, O., Rajamani, S.K.: Fair Simulation. *Inf. Comput.* **173**(1), 64–81 (2002)
22. Ilie, L., Navarro, G., Yu, S.: On NFA Reductions. In: *Theory Is Forever, Essays Dedicated to Arto Salomaa on the Occasion of His 70th Birthday*. pp. 112–124 (2004)
23. Kähler, D., Wilke, T.: Complementation, disambiguation, and determinization of büchi automata unified. In: *Proc. of ICALP’08*. pp. 724–735. Springer (2008)
24. Kupferman, O., Vardi, M.Y.: Weak Alternating Automata Are Not That Weak. *ACM Trans. Comput. Logic* **2**(3), 408–429 (2001)
25. Li, Y., Turrini, A., Zhang, L., Schewe, S.: Learning to complement büchi automata. In: *Proc. of VMCAI’18*. pp. 313–335. Springer (2018)
26. Mayr, R., Clemente, L.: Advanced automata minimization. In: *Proc. of POPL’13*. pp. 63–74 (2013)
27. Mayr, R., Clemente, L.: Efficient Reduction of Nondeterministic Automata with Application to Language Inclusion Testing. *Logical Methods in Computer Science* **15** (2019)
28. Michel, M.: Complementation is More Difficult with Automata on Infinite Words. *CNET, Paris* **15** (1988)
29. Piterman, N.: From nondeterministic büchi and streett automata to deterministic parity automata. In: *Proc. of LICS’06*. pp. 255–264. IEEE (2006)
30. Ranzato, F., Tapparo, F.: A new efficient simulation equivalence algorithm. In: *Proc. of LICS’07*. pp. 171–180 (2007)
31. Ranzato, F., Tapparo, F.: An Efficient Simulation Algorithm based on Abstract Interpretation. *Information and Computation* **208**(1), 1–22 (2010)
32. Safra, S.: On the Complexity of ω -automata. In: *Proc. of FOCS’88*. pp. 319–327. IEEE (1988)
33. Schewe, S.: Büchi complementation made tight. In: *Proc. of STACS’09*. pp. 661–672. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik (2009)
34. Sistla, A.P., Vardi, M.Y., Wolper, P.: The Complement Problem for Büchi Automata with Applications to Temporal Logic. *Theoretical Computer Science* **49**(2-3), 217–237 (1987)
35. Sistla, P.A., Vardi, M.Y., Wolper, P.: The complementation problem for Büchi automata with applications to temporal logic. In: *Automata, Languages and Programming*. pp. 465–474. Springer (1985)
36. Somenzi, F., Bloem, R.: Efficient Büchi automata from LTL formulae. In: *Proc. of CAV’00*. pp. 248–263. Springer (2000)
37. Tabakov, D., Vardi, M.Y.: Experimental Evaluation of Classical Automata Constructions. In: *Proc. of LPAR’05*. pp. 396–411. Springer (2005)
38. Vardi, M.Y.: The Büchi complementation saga. In: *Proc. of STACS’07*. pp. 12–22. Springer (2007)
39. Vardi, M.Y., Wilke, T.: Automata: From Logics to Algorithms. *Logic and Automata* **2**, 629–736 (2008)
40. Vardi, M.Y., Wilke, T., Kupferman, O., Fogarty, S.J.: Unifying Büchi Complementation Constructions. *Logical Methods in Computer Science* **9** (2013)
41. Yan, Q.: Lower bounds for complementation of ω -automata via the full automata technique. In: *Proc. of ICALP’06*. pp. 589–600. Springer (2006)